OPTIMAL INQUIRY

Tai-Wei Hu and Andriy Zapechelnyuk

Abstract. A decision maker acquires and processes information about an uncertain state of nature by an *inquiry*, which is a procedure prescribing a sequence of questions to be asked before a decision is reached. The decision maker bears a cognitive cost proportional to the length of inquiry. We characterize optimal inquiries and uncover two behavioural implications: attention span reduction (the decision maker favours shorter inquiries as cognitive cost rises by focusing on a subset of decisions and assigns them different priorities) 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: Bounded rationality; information theory; cognitive cost; rational inattention; attention span; confirmation bias; consideration set

Date: June 20, 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 Larry Samuelson and Jakub Steiner 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, with novel behavioral implications on attention span and confirmation bias.

We formalize an *inquiry* as the decision maker's strategy of asking questions about the relevant state of nature. 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 decision. Unlike the standard framework, however, our framework explicitly postulates a cognitive cost associated with the length of the inquiry.

Our framework provides an explicit and intuitive procedure for information processing. It has the same backbone motivation as what gave rise to the rational-inattention literature (surveyed in Maćkowiak et al., 2023). The main departure of our approach from this literature is that we focus on the dynamic process of inquiry with an endogenous choice of the optimal procedure. This allows us to obtain behavioral implications that are of dynamic nature, such as an endogenous preference for a shorter attention span and a prioritization of certain salient decisions before considering others.

Moreover, our cost of inquiry is directly associated with the cognitive activity of the decision maker, namely, the acts of asking questions and processing their answers, and thus it is independent of the decision maker's beliefs. For example, the cost of performing a blood-sugar test and processing its result (in terms of physical or cognitive resources) is independent of the patient's medical history. In contrast, in the standard rational inattention model, the cost is a function of the decision maker's prior beliefs, which can be unrealistic is certain applications and has a conceptual problem if applied to game situations (Denti et al., 2022).

Our main result is a characterization of optimal inquiry. To obtain this result, we overcome a technical challenge: our set of possible inquiries is an infinite set, with some of its elements (the inquiry tree) discrete and others (the information partition) continuous. This is a nonstandard problem that does not admit the standard first-order approach. Instead, we reformulate the task of finding an optimal inquiry as a simple finite optimization problem. Our characterization relies on the following two principles for optimality of an inquiry.

First, we show that an optimal inquiry is dynamically consistent in the following sense. Consider a decision maker who processes information according to an 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 acquired so far as given. Dynamic consistency requires it be optimal to stick to the original plan. We prove this property for any optimal inquiry in our framework.

Second, we use two well-known results from the information theory—the Kraft inequality (Kraft, 1949) and the Huffman coding (Huffman, 1952)—to characterize the set of payoff-relevant outcomes that can be implemented by an optimal inquiry. Any such outcome consists of two parts—*form* and *content*. The form includes a *consideration set*, which is a subset of feasible decisions that can be implemented in that outcome, and a *length profile*, which specifies how many questions are asked to reach each decision in the consideration set. The content consists of an *information partition*, which describes the posterior information about the state upon reaching each decision in the consideration set. We show that the form determines the content: given the optimal consideration set and the length profile, the boundaries of the optimal information partition are determined by simple indifference conditions.

However, the content is also informative about the form: given the information partition, the optimal length profile is determined by the Huffman coding. This also implies a negative correlation between the ex ante likelihood of choosing a decision and the inquiry length that leads to that decision. That is, more likely decisions are prioritized and considered before other decisions. To determine the optimal form, optimal inquiry trades off the accuracy of information processed, expressed as the finesse of the resulting information partition, against the cognitive resources needed to achieve it, measured by the number of questions to be asked. We draw two behavioural implications from this trade-off.

First we consider implications related to the form of optimal inquiry, and define *attention span* as the expected number of questions the DM asks before reaching a decision. We show that the DM optimally reduces her attention span as the cognitive cost rises. This is achieved either by dropping some decision out of the consideration set, or by prioritising some decisions over others, or both. At the extreme, when the cost is very low, all feasible options are considered, and it takes as many questions as needed to distinguish them all. On the other hand, when the cost is very high, no information is processed, and the decision is chosen according to the prior belief.

Second, we consider implications to the content of the optimal inquiry. We show that optimal inquiry always exhibits *confirmation bias*, in the following sense. The decision maker optimally seeks information to confirm her prevalent hypothesis of which decisions are optimal. This formalizes the informal definition of confirmation bias in psychology such as Nickerson (1998): "It refers usually to unwitting selectivity in the acquisition and use of evidence." We uncover an economic mechanism for the confirmation bias to occur optimally. Because of the cognitive cost, the decision maker is willing to make suboptimal choices that are associated with fewer questions. At the same time, ex ante more likely choices are optimally prioritized with fewer questions to confirm them. These two forces together lead to the confirmation bias endogenously.

Finally, we use our model to study two phenomena, misdiagnosis in primary health care and wrongful convictions in criminal investigation, where the literature has argued that cognitive factors are important for biases in the inquiries with dire consequences. Through the lens of our model, we show that the pressure to end inquiry early can lead to a biased process. In the case of health care, we show features such as "premature diagnosis" and "search satisficing" can be explained by our confirmation bias. In the case of criminal justice, we show that a "tunnel vision" that leads to higher rate of wrongful conviction can be linked to higher cognitive cost.

Related Literature. This paper makes a conceptual and methodological contribution to three strands of literature.

The first strand includes papers that formulate and study decision making with cognitive limitations. A prominent approach in this literature is the rational inattention approach, initiated by Sims (2003). It treats limited cognition as costly information acquisition. The cost of acquiring information is postulated as an ex-ante cost function, typically modelled as entropy reduction relative to the prior belief, as in Matějka and McKay (2015) and Jung et al. (2019). Morris and Strack (2019) introduce an alternative ex-ante cost function motivated by the classic sequential sampling problem of Wald (1945). Bloedel and Zhong (2021) provide general conditions for ex-ante cost functions to arise from dynamic models of information acquisition, and Pomatto et al. (2023) characterize ex-ante cost functions that satisfy several economically interpretable axioms.

Unlike this literature, we focus on a concrete but intuitive dynamic model where the cost of information is directly associated with asking questions. The dynamic nature of the process and the sequencing of questions matters and has behavioral implications. Moreover, our dynamic consistency result shows the decision-maker's commitment to an information acquisition strategy—which is a typical assumption in the above literature—has no power in our setting.

Cognitive limitations of a decision maker have also been modeled without reducing them to an ex-ante cost function. 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 constraint 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. Cremer et al. (2007) propose a model of organizational language using codes, with the main trade-off between the use of broader codes, which are easier to process, and the precision of such codes. While our model shares a similar trade-off, our model of inquiry is dynamic in nature with implications on the timing of information processing.

The second strand of literature includes papers that studies behavioural biases with cognitive frictions. These papers range from axiomatic to constrained optimization approaches. For consideration set, 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 of Masatlioglu et al. (2012).

The third strand rationalizes confirmation bias. The wisdom from the literature is that frictions in information processing tend to cause the decision-maker to favour signals that confirm the prior belief. Wilson's (2014) model generates this 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. Jehiel and Steiner (2020) obtain confirmation bias in a model where the decision maker chooses whether or not to continue to receive more signals, but can only remember the last one received. Confirmation bias here means that the agent is more likely to stop when seeing a signal in favour of the prior.

In the rational-inattention literature, Steiner et al. (2017) obtain a "status quo bias" in a dynamic rational-inattention model where the decision-maker tend to stick to prior decisions. Nimark and Sundaresan (2019) also obtain a "confirmation effect," meaning that the decision-maker adopts signal structures in favour of the prior belief. All these papers argue that certain implications from the proposed models can be interpreted as confirmation bias, and emphasize the importance of the prior belief. In contrast, we define confirmation bias formally as the decision-maker seeking evidence to confirm ex ante most likely guesses about which decision is optimal, a definition that is based not on priors but on observable choices.

2. The model

2.1. **Primitives.** A decision-maker (DM) needs to process information about an uncertain state of nature before taking an action. The DM's utility $u(a, x)$ depends on her action, $a \in A$, and an uncertain state, $x \in X$.[1] 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 dominates another action a' if $u(a,x) \geq u(a',x)$ for all $x \in X$ and strictly so for some $x \in X$. Throughout the paper, we assume:

(A_1) For all $a \in A$, $u(a,x)$ is continuous in x, and $\mathbb{E}[u(a,x)]$ is finite.

(A_2) For all $a, a' \in A$, a does not 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. Assumption (A_2) is introduced to simplify exposition and it 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 two classes of utility functions.

(U_1) The Lancaster model of product characteristics: $X \subset \mathbb{R}^L$ and, for each $a \in A$, there is $(\alpha_a, \beta_a) \in \mathbb{R} \times \mathbb{R}^L$ such that $u(a,x) = \beta_a \cdot x + \alpha_a$.

(U_2) A tracking problem: $A \subset \mathbb{R}^L$ and $X \subset \mathbb{R}^L$, and $u(a,x)$ is the negative distance between a and x, that is, $u(a,x) = -||a-x||_p + \alpha_a$, where $||\cdot||_p$ is the L_p-norm on $\mathbb{R}^L$ and $\alpha_a \in \mathbb{R}$ for each $a \in A$.

There are two special cases of (U_1) that we will use for illustrations. The first case has $L = 1$ and hence the utilities depend only on a one-dimensional state. The second case has $A = \{a_1, ..., a_L\}$ and $u(x, a_l) = x_l$ for $l = 1, ..., L$, where the values x_l are distributed independently. This is the case where the DM chooses between L independently valued options.

[1]Variable x can be interpreted as a profile of observables or signals with quantitative information about the true underlying state of nature (which may be ultimately unobservable) that the DM can ask questions about.

2.2. **Inquiries.** When confronted with a state x, the DM does not observe x directly. Instead, she relies on a series of questions to process information about x. Formally, we consider an *inquiry* as a series of *true/false* questions formulated as propositions. A proposition is a statement about x in the form "$x \in Y$" that can be true or false. Denote the collections of Borel subsets of X by $\mathcal{B}(X)$, and let us 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. Non-terminal nodes of the tree are associated with propositions, and terminal nodes are associated with actions. 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 for the tree 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 given the set of states X by $\mathcal{Q}_X$.

Given a state of nature $x \in X$, an inquiry $Q = \langle N, T, \sigma, \chi, d \rangle$ 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 inquiry 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 his blood sugar level is 6 *mmol/L* means little to a medical lay person, 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).

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 the 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

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

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, and we call it a *category* of states induced by Q. Note that the collection of categories $\{I_t(Q) : t \in T\}$ forms a partition of X. It is the information partition at the end of the inquiry.

As zero probability events do not matter for payoffs, 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 disjoined sets $\{X_1, X_2, ..., X_K\}$ is a *partition of* X if $\mathbb{P}(X_k) > 0$ for each k, and $\sum_k \mathbb{P}(X_k) = \mathbb{P}(X) = 1$.

Note that a partition according to the above definition does not have to be exhaustive; it is sufficient for the partition to cover a measure-one set. We adopt this definition to avoid discussions about measure-zero sets that have no rendering on the DM's expected payoffs.

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 and a state x, if the inquiry reaches the terminal node t for the given x, the DM's ex-post payoff net of the cognitive cost 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

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

We are interested in the optimal inquiry that solves

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

The maximization problem (3) 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. While from information theory 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, in our model 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, in contrast to the usual setup in which the model is silent about the corresponding procedure that the DM uses to arrive at 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 at that decision, that is, ℓ is smaller.

2.5. **Example.** We illustrate our setting and, later, the results by the following example. The example is based on the case study in Croskerry et al. (2013), which illustrates how cognitive factors affect misdiagnosis in healthcare. A detailed description and the implications from our theory to misdiagnosis will be given are Section 6.1. This example is based on the following stylized situation of a doctor visit. A

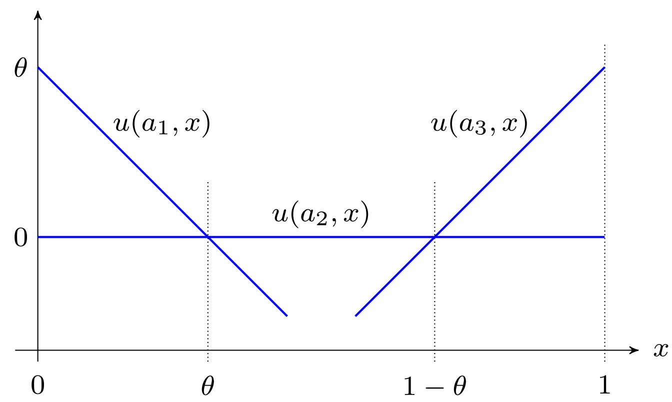

FIGURE 1. The doctor's utility from actions a_R, a_L, and a_I.

patient comes to a family doctor about a common symptom but may in fact have a rare condition that requires further investigations to avoid serious health implications. The severity of the issue is summarized by a state $x \in [0,1]$. The doctor has three possible actions: to send the patient home to rest (labeled as action a_1), to prescribe the usual medication for the common symptom (labeled as action a_2), or to refer the patient for further investigation (labeled as action a_3). Depending on the severity of the issue, x, the doctor's gross payoffs from these actions are given by the quadratic loss relative to the respective ideal states 0, 1/2, and 1:

$$U(a_1,x) = -x^2, \quad U(a_2,x) = (\tfrac{1}{4} - \theta) - (\tfrac{1}{2} - x)^2, \quad U(a_3,x) = -(1-x)^2,$$

where $\theta \in (0, 1/2)$ is a parameter capturing the importance of the extreme actions a_1 and a_3 relative to the middle action a_2. In other words, a_1 would be ideal for rather healthy patients, i.e., patients with small x's, a_2 for x's around the middle, and a_3 for severe conditions. For convenience, fix a default action, say, a_2, and consider the utility $u(x,a)$ from each action $a \in \{a_1, a_2, a_3\}$ as compared to the default action, $u(x,a) = U(x,a) - U(x,a_2)$. Thus,

$$u(a_1,x) = \theta - x, \quad u(a_2,x) = 0, \quad u(a_3,x) = \theta - (1-x),$$

as shown in Figure 1.

The doctor is initially uninformed about x. Note that the doctor does not need to discover x precisely, she only needs to find out enough to choose a treatment. To learn about x, the doctor asks several yes/no questions according to an inquiry that starts with an initial question, specifies follow-up questions depending on earlier answers,

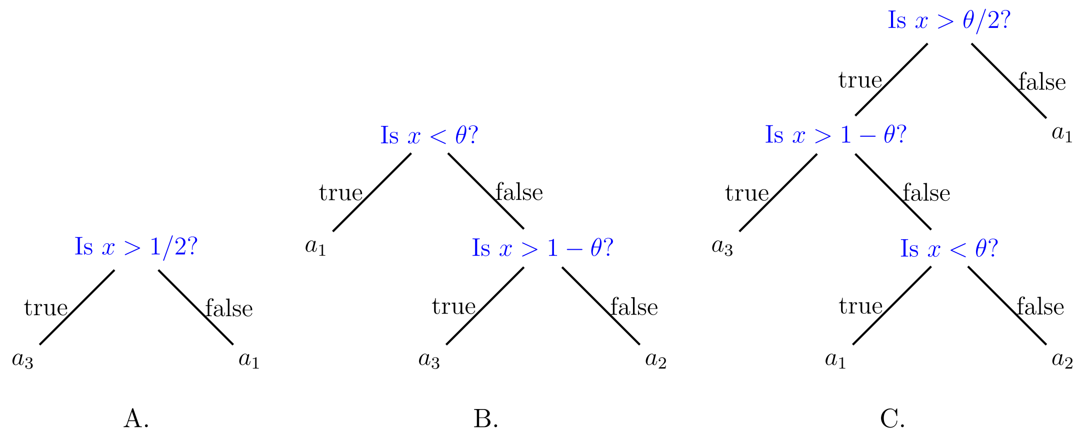

FIGURE 2. Examples of inquiries

and prescribes actions. Each question is formulated as a proposition "$x \in Y$" with $Y \subset [0, 1]$, which can be true or false. Examples of inquiries are shown in Figure 2.

A cost λ of a question is interpreted as the opportunity cost of time and cognitive effort spent on a patient that could have been spent to diagnose and treat other patients. Indeed, in Croskerry et al. (2013) this cognitive cost is regarded as an important factor that affects the doctor's investigation and the resulting decision. If the doctor reaches a decision a after asking ℓ questions, the resulting payoff is thus $u(a, x) - \lambda \ell_a$. For example, in the inquiry B (Figure 2), the cost is λ if a_1 is reached, and it is 2λ if either a_2 or a_3 are reached. The doctor would like to choose an inquiry that maximizes her expected utility net of the cost of inquiry, given his prior knowledge, modelled as a prior distribution over x.

As a benchmark, suppose that there is no cost of asking questions, $\lambda = 0$. Then, as apparent from Figure 1, it is optimal to choose a_1 when x is below θ, to choose a_2 when x is between θ and $1 - \theta$, and to choose a_3 when x is above $1 - \theta$. Inquires B and C (Figure 2) both achieve this outcome. However, when questions are costly, $\lambda > 0$, the two inquiries differ significantly in terms of the cognitive cost: when action a_1 is taken, it takes only one question in inquiry B but it may take three in inquiry C; when action a_2 is taken, it takes two questions in B but three in C. Moreover, once we take the cognitive cost into account, the doctor may find it optimal to trade off some accuracy of information about x to reduce the cost of inquiry; in other words, the optimal information partition would be endogenously determined by the cognitive cost. As we shall see later, because of these considerations inquiries B and C in Figure 2 are neither equivalent to each other nor optimal.

3. Optimal Inquiries

We first establish two principles of optimality of inquiries. We show that optimal inquiry is dynamically consistent. We also show that an inquiry can be summarized by its payoff-relevant outcome, and we then characterize the outcomes of optimal inquiries. Based on these principles, we will express the task of finding an optimal inquiry as a simple finite optimization problem, and analyze the properties of its solution.

3.1. **Dynamic Consistency.** We show that it makes no difference whether the DM commits to an optimal inquiry ex ante or she is free to update her strategy at any interim stage, and hence, the inquiry is not only ex ante, but also sequentially optimal.

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 strategy: whether to follow the original plan Q_n or to deviate to another sub-inquiry $\hat{Q} \in \mathcal{Q}_{I_n(Q)}$. Let $W_n(\hat{Q}; \lambda)$ be the DM's expected payoff conditional on reaching node n if she chooses sub-inquiry $\hat{Q}$ 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 in a sequentially optimal way at each terminal node as well. Specifically, the DM chooses an action 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.*

The theorem is proved by the typical argument for dynamic consistency. At any node n, the proof shows the following: if there is a sub-inquiry that is superior to the original one, then one can modify the original inquiry by plugging in the superior sub-inquiry after node n and obtain a strictly higher ex ante payoff. Crucial to this argument, however, is the fact that at any node n, the cognitive cost paid for the questions asked to arrive at n is sunk because of the additive-cost structure.[2]

Theorem 3.1 demonstrates the procedural rationality of the optimal inquiry, a property that cannot be discussed without the explicit formulation of the decision-making process. Moreover, as we shall see later, although the optimal inquiry features certain "biases" from the perspective of a model without the cognitive cost, these biases are not driven by inconsistent behavior between different stages of the decision process, they are an inevitable part of the optimal response to the cognitive cost.

3.2. **Outcomes.** Here we show that it suffices to describe an optimal inquiry by its payoff-relevant outcome. The outcome consists of two parts: the *form* and the *content*. The form consists of a *consideration set*—which is a subset of decisions that can be implemented in that outcome—and a *length profile*—which specifies how many questions are asked to reach each decision in the consideration set. The content consists of a collection of *categories* that forms an information partition, which describes the posterior information about the state upon reaching each decision in the consideration set.

We begin by observing 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

[2]For example, in the model where the ex-ante cost of inquiry is proportional to the longest path in the tree, at any interim node n the cost is not sunk, and, thus, dynamic consistency may not hold.

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.*

Next, we observe that an optimal inquiry cannot induce the same action in two or more terminal nodes. Indeed, if it was the case, there would have been no need to distinguish between these terminal nodes, so the number of costly questions in the inquiry could be reduced. For example, in inquiry C (Figure 2), action a_1 is chosen after a single question when $x \in [0, \theta/2]$ and after three questions when $x \in (\theta/2, \theta)$. Let us merge these conditions into a single proposition, $x < \theta$. Asking whether $x < \theta$ first, and choosing a_1 if true, and otherwise asking the remaining question, whether $x > 1 - \theta$, leads us to inquiry B. Inquiry B chooses each action on the same subset of states as inquiry C, but asks fewer questions. This observation leads us to the following property of optimal inquiry.

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*, and to actions in $D(Q)$ as *decisions*. The set $D(Q)$ can be a proper subset of A, with the interpretation that the DM will process information in a way that would lead her only to consider a strict subset of all feasible actions. For each decision $a \in D(Q)$, let $\ell_a(Q)$ denote the length of inquiry leading to the terminal node where a is chosen, and let $I_a(Q)$ denote the information set or the *category* induced by Q in that terminal node. Let $\ell(Q) = \{\ell_a(Q)\}_{a\in A}$ and $I(Q) = \{I_a(Q)\}_{a\in A}$. We will refer to the triple $(D(Q), \ell(Q), I(Q))$ as the outcome profile induced by Q, and hence $(D(Q), \ell(Q))$ describes the form of the inquiry Q and $I(Q)$ describes the content. Note that the form of an inquiry is discrete in nature while the content is continuous.

Note that actions in $A - D(Q)$ are never chosen; for every action $a \in A - D(Q)$ we use the notation $\ell_a(Q) = \varnothing$ and $I_a(Q) = \varnothing$. Since under the inquiry Q, a decision a

with a shorter inquiry length grasps the DM's attention first before another a' with a longer length, we may say that the DM *prioritizes* a decision a over another a' if $\ell_a(Q) < \ell_{a'}(Q)$.

For illustration, consider inquiries in Figure 2:

Inquiry A's outcome: $D = \{a_1, a_3\}$, $\ell = (1, \varnothing, 1)$, $I = \{[0, 1/2], \varnothing, (1/2, 1]\}$.
Inquiry B's outcome: $D = \{a_1, a_2, a_3\}$, $\ell = (1, 2, 2)$, $I = \{[0, \theta), [\theta, 1-\theta], (1-\theta, 1]\}$.

Inquiry C has the same action a_1 in two different terminal nodes, and thus cannot be represented this way. Moreover, by Lemma 3.2, inquiry C cannot be optimal. In fact, this example also shows that the subtlety needed for the proof of Lemma 3.2. Indeed, the two terminal nodes for which the same action, in this case a_1, is taken, occurs at different branches of the inquiry tree. Thus, to show that the inquiry C is suboptimal, we construct another inquiry (namely, inquiry B) which leads to the same information partition but with the categories for the two terminal nodes with action a_1 combined into one single category, and with all the inquiry shorter than the corresponding ones in inquiry C.

This construction is based on a more general principle that leads to the following characterization of inquiry outcomes. Let $D \subseteq A$, let $\ell = (\ell_a)_{a\in A} \in (\mathbb{N} \cup \{\varnothing\})^{|A|}$ be a length profile, and let I be a partition of X. Denote by $\mathcal{Z}$ the set of such triples (D, ℓ, I). We say that an outcome profile $(D, \ell, I) \in \mathcal{Z}$ is *implementable* if there exists an inquiry $Q \in \mathcal{Q}_X$ that induces this outcome profile, that is, $(D, \ell, I) = (D(Q), \ell(Q), I(Q))$. The following lemma characterizes implementable outcomes.

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

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

Equality (6) follows from the Kraft inequality in information theory that characterizes the path lengths of binary trees. Here we have equality instead of inequality because in our inquiry trees each non-terminal node has precisely two outgoing branches.

Lemma 3.3 implies that the set of feasible outcomes only depends on the form of an inquiry, that is, the consideration set D and the length profile ℓ, but it does not depend on the content, that is, the information partition. In other words, the lemma shows that for any given form (D, ℓ) that satisfies (6) and any given content I with $|D|$ non-empty categories, we can construct an inquiry with the corresponding outcome. However, as we shall see later, not every information partition is optimal, and that the form and the content are interdependent.

To illustrate this lemma, let us return to our example in Section 2.5. There are only three actions in the example. Thus, only three length profiles for $D = \{a_1, a_2, a_3\}$ satisfy equality (6), namely, $\ell = (1, 2, 2)$ (see inquiry B in Figure 2), $\ell = (2, 1, 2)$, and $\ell = (2, 2, 1)$. If we increase the number of actions to four, so $D = \{a_1, a_2, a_3, a_4\}$, then there will be 13 length profiles that satisfy (6), namely, the uniform profile, $(2, 2, 2, 2)$, and 12 distinct permutations of the extreme profile $(1, 2, 3, 3)$. This set of length profiles increases exponentially with the size of D. However, later on we develop optimal conditions with which we can identify a smaller candidate set.

Lemmas 3.1–3.2 lead us to a key observation. An outcome (D, ℓ, I) captures all we need to know to evaluate the DM's expected payoff of an inquiry that leads to that outcome. Indeed, suppose that two different inquiries Q and Q' implement the same outcome (D, ℓ, I). Then, by (2), we have

$$W(Q; \lambda) = W(Q'; \lambda) = \sum_{d \in D} \int_{x \in I_d} \Big(u(d, x) - \lambda \ell_d \Big) G(\mathrm{d}x). \tag{7}$$

Moreover, Lemma 3.3 implies that for any outcome (D, ℓ, I) that satisfies (6) there exists an inquiry with that outcome.

Thus, without loss of generality, an inquiry can be equivalently represented by its outcome (D, ℓ, I). An outcome of an optimal inquiry will be called *optimal outcome*.

3.3. **Optimal Inquiries.** We have shown an inquiry can be summarized by its outcome (D, ℓ, I). Moreover, Lemma 3.3 shows that the information partition I does not affect whether or not an outcome profile if implementable. This characterization allows us to solve the optimal inquiry problem in two stages. We first fix an arbitrary form (D, ℓ) that satisfies (6), and solve for the optimal content $I = I^*(D, \ell)$. Then, we maximize over all possible forms of (D, ℓ)'s.

In the first stage, taking (D, ℓ) as given, we find an information partition $I^*(D, \ell)$ that maximizes the DM's expected utility. Specifically, let $I^*(D, \ell) = \{I^*_d(D, \ell)\}_{d \in A}$,

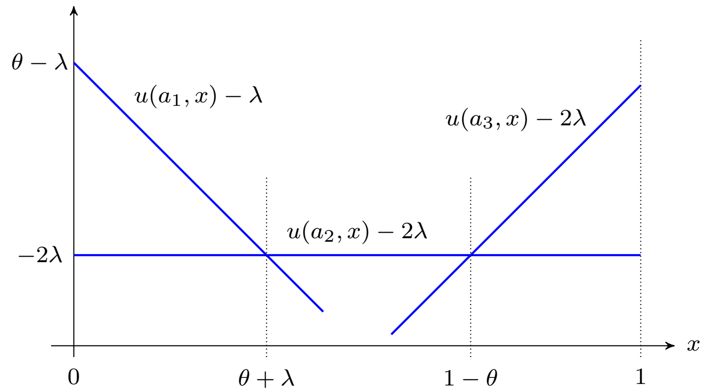

FIGURE 3. Determination of I^* from $\ell = (1, 2, 2)$

where

$$I^*_d(D, \ell) = \begin{cases} \left\{x \in X : u(d, x) - \lambda \ell_d > \max_{a \in D - \{d\}} u(a, x) - \lambda \ell_a \right\}, & d \in D, \\ \varnothing, & d \in A - D. \end{cases} \tag{8}$$

That is, for each decision d that can be chosen, so $d \in D$, $I^*_d(D, \ell)$ is 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. Note that $I^*(D, \ell)$ is a partition of X according to Definition 2.1, because, by assumption (A_3), the set $\left(X - \bigcup_{d \in D} I^*_d(D, \ell)\right)$ has measure zero.

We have the following lemma.

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

The key observation to understand this lemma is that $I^*(D, \ell)$ is the optimal information partition given the form, (D, ℓ), as the DM chooses the unique best-response action for each state $x \in X$, except for a measure zero of states. Indeed, given (D, ℓ), the "effective" utility for the DM by choosing action d is now $u(d, x) - \lambda \ell_d$, when taking the cognitive cost into account.

Lemma 3.4 is essential to solving the optimal inquiry. As mentioned earlier, the outcomes of an inquiry include both a continuous element I and discrete element (D, ℓ). Lemma 3.4 shows that of the optimal content I is determined by the form (D, ℓ) through $I^*(D, \ell)$. It also generates candidate optimal inquiries effectively. For example, it immediately implies that Inquiry B in Figure 2 is suboptimal. Indeed, Inquiry B has the partition

$$I_{a_1} = [0, \theta),\ I_{a_2} = [\theta, 1-\theta],\ I_{a_3} = (1-\theta, 1],$$

as shown in Figure 1. But, given the length profile is $\ell = (1, 2, 2)$ with the associated cost λ of each question, the DM can do strictly better by using the partition

$$I^*_{a_1} = (0, \theta+\lambda),\ I^*_{a_2} = (\theta+\lambda, 1-\theta),\ I^*_{a_3} = (1-\theta, 1), \tag{9}$$

as shown in Figure 3. Indeed, on the interval $(\theta, \theta+\lambda)$, the DM optimally chooses decision a_1, even though a_2 would have been a better decision absent cognitive cost. This is because a_1 needs one less question to ask and, thus, saves λ, while $u(x, a_2) - u(x, a_1) < \lambda$ for any state x in that interval.

As a consequence of Lemma 3.4, the maximization problem (3) can now be reduced to the choice of the form, (D, ℓ). Let $\mathcal{F}^*$ be the set of all forms (D, ℓ) with $D \subseteq A$ and ℓ satisfying (6). The DM chooses a form $(D, \ell) \in \mathcal{F}^*$, 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*

$$\max_{(D,\ell)\in\mathcal{F}^*} \sum_{d\in D} \int_{x\in I^*_d(D,\ell)} \Big(u(d, x) - \lambda \ell_d\Big) G(\mathrm{d}x). \tag{10}$$

Because $\mathcal{F}^*$ 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.*

Another straightforward implication of Theorem 3.2 is that the optimal inquiry satisfies a minimal rationality property: the independence of irrelevant alternatives principle. This property is also known in the literature as "attention filter" (Masatlioglu et al., 2012). It is defined as follows. Suppose that DM's consideration set D is a strict subset of A. Then, the attention filter property requires that, for any a smaller action set $A' \subset A$ that contains D, the optimal consideration set is still D. This property holds under optimal inquiry.

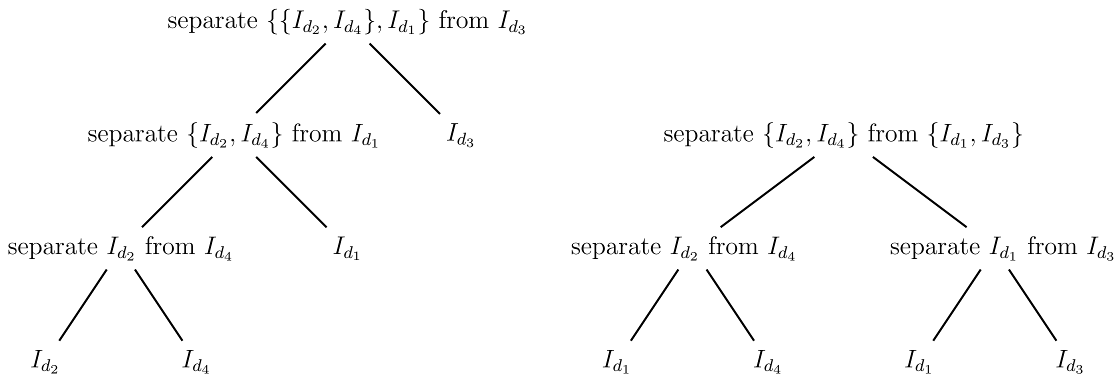

(a) $\ell = (2, 3, 1, 3)$ (b) $\ell = (2, 2, 2, 2)$

FIGURE 4. Huffman coding with $|D| = 4$ and $\mathbb{P}(I_{d_3}) > \mathbb{P}(I_{d_1}) > \mathbb{P}(I_{d_2}) > \mathbb{P}(I_{d_4})$

Corollary 3.2. *If (D, ℓ, I) is an optimal outcome for action set A, then it is also an optimal outcome for each action set A' such that $D \subseteq A' \subsetneq A$.*

Although $\mathcal{F}^*$ is a finite set, as mentioned earlier, it can be a relatively large set. However, there is an additional optimality condition that helps determine the optimal length profile more efficiently, and it also helps characterize its behaviour. Indeed, while Lemma 3.4 characterizes the optimal categories for a given length profile, one can also look for optimality conditions for the length profile for given categories. Specifically, given a partition $\{I_d\}_{d\in D}$, the optimal ℓ must minimize the average length with respect to the distribution such that ℓ_d occurring with probability $\mathbb{P}(I_d)$ and subject to the constraint (6). This is a well-known problem in information theory, and the solution is described by the algorithm called *Huffman coding*. Here we show how the algorithm works for $|D| = 4$. The generalization to arbitrary D is straightforward. We refer to, e.g., Thomas and Cover (2006), Section 5.6, for formal details.

Consider $D = \{d_1, d_2, d_3, d_4\}$ with the following probability ranking of the decisions:

$$\mathbb{P}(I_{d_3}) > \mathbb{P}(I_{d_1}) > \mathbb{P}(I_{d_2}) > \mathbb{P}(I_{d_4}). \tag{11}$$

Specifically, let

$$\mathbb{P}(I_{d_1}) = 0.25,\ \mathbb{P}(I_{d_2}) = 0.2,\ \mathbb{P}(I_{d_3}) = 0.4,\ \mathbb{P}(I_{d_4}) = 0.15.$$

In stage $t = 0$, let us define $p^0_d = \mathbb{P}(I_d)$ for each $d \in D$, and order the decisions according to their probabilities: $p^0_{d_3} = 0.4 > p^0_{d_1} = 0.25 > p^0_{d_2} = 0.2 > p^0_{d_4} = 0.15.$

We then merge the last two, $\{d_2, d_4\}$, and treat the pair as a single decision whose probability is $p^1_{\{d_2,d_4\}} = 0.2 + 0.15 = 0.35$. The other decisions and their probabilities stay the same: $p^1_{d_1} = p^0_{d_1}$ and $p^1_{d_3} = p^0_{d_3}$. In stage $t = 1$, we reorder the decisions of stage $t = 0$ according to their probabilities: $p^1_{\{d_3\}} = 0.4 > p^1_{\{d_2,d_4\}} = 0.35 > p^1_{\{d_1\}} = 0.25$. We then merge the last two, $\{\{d_2, d_4\}, \{d_1\}\}$, and treat the set as a single decision with probability $p^2_{\{\{d_2,d_4\},\{d_1\}\}} = 0.35 + 0.25 = 0.6$. In stage $t = 2$, again, we reorder the decisions of stage $t = 1$: $p^2_{\{\{d_2,d_4\},\{d_1\}\}} = 0.6 > p^2_{\{d_3\}} = 0.4$. We then merge the remaining decisions to obtain $\{\{\{d_2, d_4\}, \{d_1\}\}, \{d_3\}\}$. Finally, we construct the inquiry tree by unraveling the nested set $\{\{\{d_2, d_4\}, \{d_1\}\}, \{d_3\}\}$ from the top layer down, as shown in Figure 4(a). The length profile for this tree is $\ell = (2, 3, 1, 3)$.

If we consider the same initial probability ranking of the decisions (11) but different probabilities,

$$\mathbb{P}(I_{d_1}) = 0.25,\ \mathbb{P}(I_{d_2}) = 0.2,\ \mathbb{P}(I_{d_3}) = 0.37,\ \mathbb{P}(I_{d_4}) = 0.18,$$

then the Huffman coding procedure yields a different inquiry tree, with length profile $\ell = (2, 2, 2, 2)$, as shown in Figure 4(b). In fact, as follows from the next proposition, $\ell = (2, 3, 1, 3)$ and $\ell = (2, 2, 2, 2)$ are the only length profiles that can be obtained given the probability ranking (11).

We now show that for any candidate consideration set D, the optimal length profile ℓ is determined by the partition I. Moreover, decisions that take longer to reach are less likely to be chosen.

Proposition 3.1. *If (D, ℓ, I) is an optimal outcome, 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 $\mathbb{P}(I_d) > \mathbb{P}(I_{d'})$, then $\ell_d \leq \ell_{d'}$.

Proposition 3.1 (b) highlights a negative correlation between the ex ante probability of a decision and the inquiry length to reach that decision. It allows us to simplify the solution of problem (10), but reducing the set of candidate outcomes $(D, \ell, I^*(D, \ell))$. Indeed, as illustrated by the above example with $D = \{d_1, d_2, d_3, d_4\}$ and probability ranking of decisions (11), out of 13 feasible length profiles that satisfy (6), only two are consistent with (11), namely, $\ell = (2, 3, 1, 3)$ and $\ell = (2, 2, 2, 2)$.

Proposition 3.1(a) is particularly useful to determine optimal ℓ for λ small. Indeed, for λ sufficiently small, the optimal length profile is determined by the information

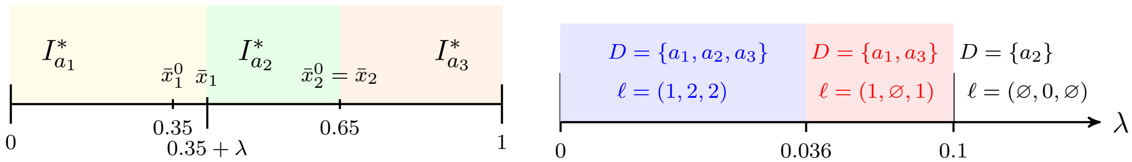

(a) Optimal information partitions

(b) Optimal (D, ℓ) as λ increases

FIGURE 5. Optimal inquiry for $\theta = 0.35$ as λ increases

partition $I^0 = \{I^0_a\}_{a\in A}$ that is optimal under standard Bayesian analysis with zero cost, $\lambda = 0$. Specifically, when the DM learns the state x for free, she simply chooses the best action for each $x \in X$. That is, for each $a \in A$,

$$I^0_a = \left\{ x \in X : u(x, d) > \max_{a' \in A - \{a\}} u(x, a') \right\}.$$

As λ increases, the information partition is continuously adjusted according to $I^*(D, \ell)$ given by (8). However, as long as λ is small enough, the optimal consideration set remains $D = A$, and the optimal length profile remains the same as the one determined by the Huffman coding for $\lambda = 0$.

For illustration, we return to the example in Section 2.5 with $A = \{a_1, a_2, a_3\}$ and utility functions given by Figure 1, and the parameter $\theta = 0.35$. When $\lambda = 0$, we have $D = A$ and $\mathbb{P}(I^0_{a_1}) = \mathbb{P}(I^0_{a_3}) = 0.35$ and $\mathbb{P}(I^0_{a_3}) = 0.3$. The optimal information partition for $\lambda = 0$ is shown in Figure 5(a), where $\bar{x}^0_1$ and $\bar{x}^0_2$ denote the thresholds between the partition elements. Using the Huffman coding, we obtain length profiles $(1, 2, 2)$ and, by symmetry, $(2, 2, 1)$. Let us fix $\ell = (2, 2, 1)$.

When $\lambda \in (0, 0.036)$, the same (D, ℓ) remain optimal, but optimal information partition $I^*(D, \ell)$ is adjusted to take into account the cost, namely, it is given by (9). The optimal information partition for such λ is shown in Figure 5(a), where $\bar{x}_1$ and $\bar{x}_2$ denote the thresholds between the partition elements. It can be seen that the category $I^*_{a_1}$ where a_1 is chosen expands as λ increases. This expansion reflects the general principle to resolve the key trade-off in our setting. On the one hand, the DM likes the categories to match with the benchmark partition I^0 to achieve higher utilities. On the other hand, the DM likes to have shorter inquiry lengths in expected terms. As λ increases, the latter becomes more important, and the DM optimally adjusts her categories to shorten the average inquiry length. Indeed, as $I^*_{a_1}$

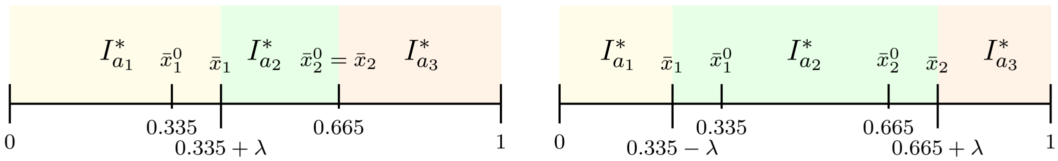

(a) $\lambda < 0.01$ and $\ell = (1, 2, 2)$

(b) $\lambda > 0.01$ and $\ell = (2, 1, 2)$

FIGURE 6. Optimal information partitions for $D = \{a_1, a_2, a_3\}$ and $\theta = 0.335$.

is associated with length 1 and others with length 2, expanding $I^*_{a_1}$ leads to shorter expected inquiry length.

For λ higher we have a similar effect but at the extensive margin by changing the consideration set. That is, for $\lambda \in (0.036, 0.1)$, the optimal D becomes $\{a_1, a_3\}$, and a_2 is no longer considered. As a result, at $\lambda = 0.036$, the category $I^*_{a_1}$ expands discontinuously, from $(0, 0.35 + \lambda)$ to $(0, 0.5)$, with a discrete jump of average inquiry length to a lower level, because of the change of the optimal consideration set.

Lastly, for $\lambda > 0.1$, it is optimal for the DM not to ask any questions, and simply choose the ex-ante optimal decision a_2. Figure 5(b) shows how the optimal pair (D, ℓ) changes as λ increases.

Curiously, there can also be changes to optimal information partition due to changes in the optimal length profile alone, while the consideration set remains the same. To illustrate this, consider the same example, but now with $\theta = 0.335$. In this case, for $\lambda < 0.043$, optimal $D = A$. As before, for λ very small, an optimal length profile is $(1, 2, 2)$, and the optimal information partition is shown in Figure 6(a). However, in this case for $\lambda > 0.01$, the optimal length profile changes to $(2, 1, 2)$, and the optimal information partition becomes as shown in Figure 6(b). This curious switch illustrates the aforementioned principle that the DM likes to shorten average inquiry as the cognitive cost rises. Indeed, as λ increases, the expected inquiry length of the length profile $(2, 1, 2)$ with the associated category I^* decreases faster than that of $(1, 2, 2)$, because with $(2, 1, 2)$ the expansion of $I^*_{a_2}$ happens on both sides while with $(1, 2, 2)$ on one side. As a result, the increasing importance of shorter inquiry length prompts the DM to switch to $(2, 1, 2)$ for λ's sufficiently high.

These two examples show that, as the cognitive cost rises, the DM respond by either expand the category with the shortest inquiry length (as in the left panel of Figure 5), or by dropping some decisions out of consideration altogether (as in the left panel of Figure 5 for higher λ's). Note that, consistent with Proposition 3.1, in the former case the expanded category is also the one with higher probability, and hence this expansion leads to lower average inquiry length, so as the latter case. In the next section we discuss this effect more generally.

4. Attention Span

Our model of inquiry can be interpreted as an attention strategy, whereby the DM focuses on various decisions during her inquiry process. With this interpretation, a natural question is then how cognitive cost affects the DM's attention span, defined as how long she would concentrate on the task of gathering information before making a decision. Formally, we measure *attention span* in our framework as the *expected inquiry length* given by

$$\bar{\ell}(D, \ell, I) = \sum_{d \in D} \ell_d \mathbb{P}(I_d). \tag{12}$$

Importantly for our purpose, it captures whether there is a lot of probability weight on a few decisions with short inquiry length, or whether this weight is more spread out among many decisions. A smaller $\bar{\ell}(D, \ell, I)$ means a shorter attention span.

In the extreme, the DM has no attention span at all when she chooses a single decision without asking any questions, in which case we have $\bar{\ell}(D, \ell, I) = 0$. The opposite extreme occurs when the lengths are all equal and the probabilities spread out. Given an outcome (D, ℓ, I), we say that the length profile ℓ is *uniform* if it assigns the same length to all decisions in D, so $\ell_d = \ell_{d'}$ for all $d, d' \in D$. In other words, the inquiry outcome is uniform if the same number of questions is asked for all states of nature. Note that this can only happen if $|D| = 2^k$ for some $k \in \{0, 1, ...\}$. Since under $|D| = 2^k$ it is always feasible to set all lengths equal to k, k is also an upper bound for the expected inquiry length for an optimal inquiry. In other words, in any optimal inquiry with $|D| = 2^k$, we have $\bar{\ell}(D, \ell, I) \leq k$, and this upper bound is achieved if and only if the length profile is uniform.

Our key result in this section is that, under optimal inquiry, a higher cognitive cost shortens the DM's attention span. Specifically, we show that the expected inquiry length decreases with λ, and strictly so when the optimal inquiry is non-uniform. Since the expected length may not be differentiable w.r.t. λ, we use the following notion of strict monotonicity. A function is strictly increasing (decreasing) *locally* at λ if there exists an interval $[\lambda_1, \lambda_2]$ with $\lambda_1 < \lambda_2$ and $\lambda \in [\lambda_1, \lambda_2]$ such that the function is strictly increasing (decreasing) on $[\lambda_1, \lambda_2]$.

Theorem 4.1. *Given λ, let $(D^\lambda, \ell^\lambda, I^\lambda)$ be an optimal outcome. The average inquiry length $\bar{\ell}(D^\lambda, \ell^\lambda, I^\lambda)$ is decreasing in λ. Moreover, it is strictly decreasing locally at λ whenever ℓ^λ is not uniform.*

Theorem 4.1 shows that higher cognitive cost always shortens the attention span, and strictly so as long as the optimal inquiry length is not uniform. The intuition for Theorem 4.1 is based on the following trade-off that the optimal inquiry resolves. On the one hand, to achieve a high (expected) utility from actions, it needs to minimize the mismatch between its category for the action and the set of states for which the action is ex post optimal; on the other hand, it needs to minimize the expected length of inquiry. As the cognitive cost increases, the latter motive becomes more important, and optimal inquiry shifts probabilities toward categories with shorter inquiries at the expense of more mismatches.

This preference for shorter inquiries generates a "bias" if we compare the information partition thus generated to the ones that would be used by a Bayesian DM under zero cognitive cost. We call this effect *confirmation bias*. The bias is endogenously determined by the cost and the utility function. As shown in Figure 6, in case of $\lambda < 0.01$ the bias favors action a_1 by expanding the set of states where a_1 is chosen (at the detriment of a_2), but for $\lambda > 0.01$ it favors action a_2 (at the detriment of both a_1 and a_3).

As mentioned earlier, this bias is generated by the motive to decrease the expected inquiry length, and this can be achieved by adjusting the inquiry either through the form or through the content. The content affects the intensive margin, and the DM can increase the probability of choosing decisions with shorter inquiry length. The form affects the extensive margin, and the DM can simply drop certain actions from the consideration set and in this way the overall inquiry length may be reduced. The former factor is addressed in detail in the next section. In the next subsection, we analyze the latter, namely how optimal consideration set is determined.

4.1. **Optimal Consideration Sets.** In our setting, a consideration set can be regarded as a set of actions that the DM deems "cognitively viable", and any action outside this set is simply ignored in the decision-making process, even though it might be ex post optimal. One important factor that determines which actions are viable is the underlying preferences. Specifically, it is useful to distinguish two actions only if they produce sufficiently different payoffs in 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:

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

The following result shows that a consideration set will optimally drop one of the two similar actions.

Proposition 4.1. *If actions a and a' are such that $\delta(a', a'') < \lambda$, then at most one of them will be in the optimal consideration set.*

Next, we show that when the cognitive cost is small enough, then all actions are optimally considered, and when the cognitive cost is large enough, then the DM asks no questions and chooses the same action in all states of nature.

Proposition 4.2. *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 $\bar{\ell}(Z_\lambda) = 0$.*

It is tempting to generalize Proposition 4.2 by conjecturing that, as λ increases, the optimal consideration set monotonically shrinks in the set inclusion order. However, this is not true in general. To illustrate this, consider the example in Section 2.5 with $A = \{a_1, a_2, a_3\}$, parameter $\theta \in (0, 1/2)$ that captures the preference for extreme actions a_1 and a_3 relative to middle action a_2, and utilities given by Figure 1.

Figure 7 shows how the optimal consideration set depends on the cost λ and the preference parameter θ. We point out three features of the optimal inquiry that may be of interest. First, the optimal consideration sets can be disjoint: for a fixed λ that is relatively high (say, for $\lambda = 0.1$), as θ increases, the optimal D changes from $\{a_2\}$ to $\{a_1, a_3\}$. In words, preferences can affect the "cognitive viability" faced by the DM, which can change in a discontinuous way. Second, the optimal consideration sets can change by multiple actions at a time: for a fixed θ that is relatively low (say, for $\theta = 0.2$), as λ increases, the optimal D changes from $\{a_1, a_2, a_3\}$ to $\{a_2\}$. Thus,

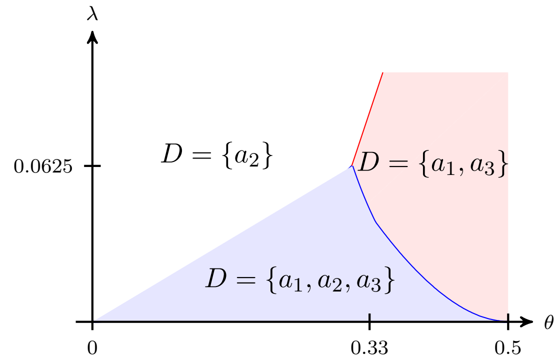

FIGURE 7. Optimal consideration sets

for two DM's with exactly the same preference, a small difference of cognitive cost can make one DM to consider all actions while making the other to consider only one. Third, an option can be phased out and then phased back in. For a given θ that is close to 1/3, say, for $\theta = 0.335$, the optimal consideration set changes from $\{a_1, a_2, a_3\}$ to $\{a_1, a_3\}$ to $\{a_2\}$, so action a_2 is dropped, but then reintroduced.

4.2. **Example with Independent Values.** In the above example where the state is of one-dimensional, there is no natural sense of how to rank the actions under assumption (A_2). However, there is a natural ranking in environments where the values of the actions are independently distributed. Here we study how this ranking affects the optimal consideration set.

Consider the model where $X = \mathbb{R}^L$ and $A = \{a_1, ..., a_L\}$, with

$$u(a_l, x) = x_l \text{ for all } l = 1, ..., L. \tag{13}$$

Assume that the values $x_1, ..., x_L$ are independently distributed. Specifically, each x_l has a distribution G_l, and $G(x) = \Pi_{l=1}^{L} G_l(x_l)$. We have the following result.

Proposition 4.3. *Suppose that $G_1 \succ_{FOSD} G_2 \succ_{FOSD} \cdots \succ_{FOSD} G_L$. Then there exists $K \in \{1, ..., L\}$ such that $D = \{a_1, ..., a_K\}$ is the optimal consideration set.*

According to Proposition 4.3, when the actions are ranked by the first-order stochastic dominance, an action can be in the optimal consideration set only if all the higher-ranked actions are in there. Moreover, if an action a_l is not in the consideration set, the DM will ask no questions about its value (x_l), nor about values of all

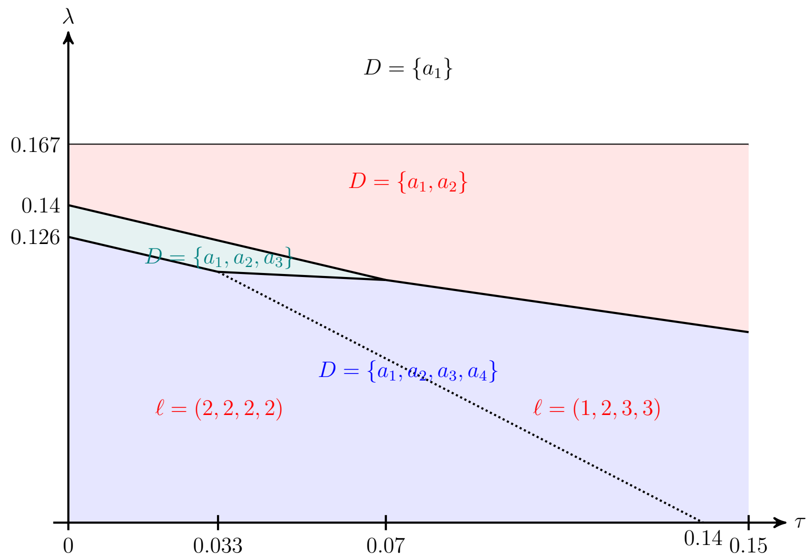

FIGURE 8. Optimal consideration set and length profile

the lower-ranked actions. That is, the DM will only spend cognitive resources on the dimensions that she deems most likely to be optimal from the ex ante perspective.

To illustrate Proposition 4.3, consider the following example. Let $L = 4$ and $A = \{a_1, a_2, a_3, a_4\}$, and the utility is given by (13). Let $\tau \in [0, 0.15]$ be a parameter. Suppose that the values x_1 and x_2 are each uniformly distributed on $[0, 1]$, while x_3 and x_4 are each uniformly distributed on $[-\tau, 1-\tau]$. Clearly, x_1 and x_2 first-order stochastically dominate x_3 and x_4.

Figure 8 shows the optimal consideration sets and length profiles (up to the symmetry between a_1 and a_2 and between a_3 and a_4) for different values of τ and λ. Note that when $|D| = 1$ and $|D| = 2$ are optimal (the white and the red areas, respectively), the expected values do not vary with τ, and hence the boundary between the white and the red areas is a horizontal line. The area where $|D| = 3$ is optimal (the green area) appears only if $\tau < 0.07$, and the corresponding optimal length profile is $(1, 2, 2)$. In all those areas, as λ increases, the adjustment comes from the extensive margin, where the number of actions considered decreases from 3 to 2, to 1. This may be regarded as focused consideration: the DM only asks questions about the first two or three values and decides based on this evidence, but ignores any potential evidence from other dimensions (e.g., x_4).

In contrast, in the area where $|D| = 4$ is optimal (the blue area), there is a shift in the length profile. For small τ (to the left of the dotted line), the optimal length profile is the uniform one, $(2, 2, 2, 2)$, as the probabilities of each action being the optimal one are not too different from one another. However, for large τ (to the right of the dotted line), the optimal length profile will switch from the uniform one to the extreme one, $(1, 2, 3, 3)$.

Thus, although the optimal inquiry does not adjust the consideration set D as λ increases, it does change the inquiry strategy discontinuously: for higher λ's, the action a_1 has the shortest length, and that set of states where a_1 is chosen will expand with λ. The latter may be interpreted as a confirmation bias in terms of the content of the inquiry: among those actions the DM is willing to consider, she is happy to expand the evidence to admit a certain action as acceptable (in this case, a_1), which would be her most likely action ex ante. In the next section we show that this is a prevalent feature of the optimal inquiry.

5. Confirmation Bias

We have seen from the previous section that, under optimal inquiry, the DM often prioritizes some actions by asking questions that lead to these decisions first, and turning to other actions only if the initial answers are negative. Moreover, some actions may not be considered at all. This may be interpreted as a form of confirmation bias in the *extensive margin*, as the DM searches for evidence to support higher-priority actions, and does not attempt to find evidence in support of actions outside the chosen consideration set. This is mainly related to the form of the inquiry.

In this section, we turn to confirmation bias in terms of the *content*, taking the form of the inquiry as given. We will show that, given the form, (D, ℓ), the DM optimally expands the categories associated with the more likely actions, relative to the zero cognitive-cost benchmark. This can be interpreted as the DM searching for evidence to confirm the desirability of the actions in D that are most likely to be optimal.

To define confirmation bias, let us consider the zero-cost case as a benchmark, and compare the set of states under which the most likely actions are taken under the optimal inquiry with and without the cognitive cost. To do so, we first rank the actions according to their likelihood under the optimal inquiry. For a fix $\lambda > 0$, let (D, ℓ, I) be an optimal outcome under cost λ, and let $K = |D|$. We order the actions in D according to how likely they are chosen under optimal inquiry, so $D = \{d_k\}_{k=1}^K$, such that

$$\mathbb{P}(I_{d_1}) \geq \mathbb{P}(I_{d_2}) \geq ... \geq \mathbb{P}(I_{d_K}), \\ \text{with a tie-breaking rule } \ \mathbb{P}(I_{d_k}) = \mathbb{P}(I_{d_{k+1}}) \implies \ell_{d_k} \leq \ell_{d_{k+1}}. \tag{14}$$

For each $k = 1, ..., K-1$, let $D_k = \{d_1, ..., d_k\}$ be the subset of k most likely actions in D. Let E_k^λ be the event that some action in D_k is preferred to all actions in $D - D_k$, when the cost λ is taken into account:

$$E_k^\lambda = \left\{ x \in X : \max_{k'=1,...,k} u(d_{k'}, x) - \lambda \ell_{d_{k'}} > \max_{m=k+1,...,K} u(d_m, x) - \lambda \ell_{d_m} \right\}, \tag{15}$$

It can be easily seen that E_k^λ coincides with $\bigcup_{d \in D_k} I_d^*(D, \ell)$ except, possibly, on a measure zero set. In words, conditional on event E_k^λ, the optimal inquiry almost surely leads to an action in D_k. Similarly, let E_k^0 be the event of choosing one of the actions from D_k in the zero-cost benchmark.

Definition 5.1. An inquiry Q with outcome (D, ℓ, I) has *confirmation bias* if for every order $\{d_k\}_{k=1}^K$ that satisfies (14),

$$E_k^0 \subseteq E_k^\lambda \text{ for all } k = 1, 2, ..., K-1. \tag{16}$$

It has *strict confirmation bias* if (16) holds, and there exists $k \in \{1, ..., K-1\}$ such that

$$E_k^\lambda - E_k^0 \text{ has a non-empty interior.} \tag{17}$$

In words, the DM has confirmation bias when affected by the cognitive cost if, for each $k = 1, ..., K-1$, she confirms to k most likely actions: she chooses one of these actions on a larger set of states, as compared to what she would have done without cognitive cost. The difference $E_k^\lambda - E_k^0$ is the set of states where an error relative to the zero-cost benchmark occurs. In any state that belongs to $E_k^\lambda - E_k^0$, the DM is biased, as she chooses an action that is suboptimal from the perspective of the Bayesian DM who knows the state.

This definition formalizes the notion of confirmation bias usually adopted in psychology. According to Nickerson (1998), "it refers usually to unwitting selectivity in the acquisition and use of evidence," which he believes is also the definition used by general psychologists. Our definition has the advantage of formally defining both confirming sets and biases: the actions the DM confirms to are the most likely ones in her optimal strategy, and the errors are defined against the zero-cost benchmark.

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

An immediate implication of Theorem 5.1 is that the probability distribution over actions over $D = \{d_k\}_{k=1}^K$ induced by the optimal inquiry first order stochastically dominates that induced by the zero-cost benchmark:

$$\mathbb{P}(E_k^\lambda) \geq \mathbb{P}(E_k^0), \quad \text{for each } k = 1, ..., K-1, \tag{18}$$

and the inequality is strict for some k if the confirmation bias is strict.

Theorem 5.1 shows that the optimal inquiry always features confirmation bias, and it has strict confirmation bias when its length profile is not uniform. From Proposition 3.1(b) we know that the most likely actions are associated with shorter inquiry lengths. At the same time, given the optimal length profile, the optimal information partition given by (8) has the feature that decisions associated with shorter inquiry lengths will be chosen on a larger set of states relative to the no-cost benchmark. These two factors reinforce each other and give rise to the confirmation bias in our setting.

We have define confirmation bias as a comparison against the benchmark case without the cognitive cost. Now we show that, this bias grows as the cognitive cost increases, in the sense that likelihood of choosing one of k ex ante most likely actions under optimal inquiry only increases as λ goes up.

Definition 5.2. Confirmation bias is *locally increasing at* λ if $\mathbb{P}(E_k^\lambda)$ is locally increasing at λ for each $k = 1, ..., K-1$. Moreover, confirmation bias is *locally strictly increasing in* λ if $\mathbb{P}(E_k^\lambda)$ is locally strictly increasing in λ for some k.

We have the following result.

Proposition 5.1. *The confirmation bias is locally increasing in λ. Moreover, it is strictly increasing locally at λ if and only if the optimal length profile at λ is not uniform.*

Proposition 5.1 shows that as the cognitive cost rises, the DM would optimally make more "errors" and is biased more toward the most likely actions. This result is

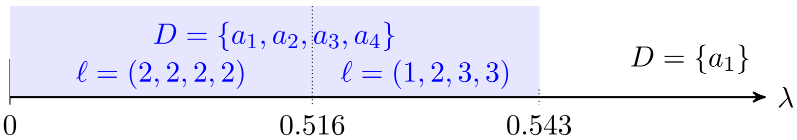

FIGURE 9. Optimal consideration set and length profile under i.i.d. exponential distribution

closely related to Theorem 4.1, which states that the DM has more "focus" on most likely actions. Naturally, as the cognitive cost rises, the DM prefers to shorten the inquiry length, and Proposition 5.1 shows that the way to optimally achieve that is by decreasing the accuracy of her categories in favour of most likely actions, which are also actions associated with the shortest inquiry lengths.

This endogenous emergence of confirmation bias in our model gives novel implications to behaviour as a result of cognitive limitations. Here we give an example that illustrates this novelty and the difference from the literature on rational inattention. Matějka and McKay (2015) show that symmetric actions will be treated symmetrically in the rational inattention model. In contrast, in our model the DM may optimally treat symmetric actions asymmetrically to save the cognitive cost. Consider the model of independent values, as in Section 4.2, with $A = \{a_1, a_2, a_3, a_4\}$, $X = \mathbb{R}^4_+$, and each x_l has exponential distribution $G_l(x_l) = 1 - e^{-x_l}$. Figure 9 depicts the optimal consideration set and the optimal length profile as a function of λ, up to the symmetry between the actions. As apparent from the figure, once λ increases above 0.516 so it becomes too costly to differentiate all actions and treat them equally, the DM prioritizes an arbitrary action (in this case, a_1) and increases the set of states where this action is chosen, eventually choosing this action alone as λ increases above 0.543. This implies endogenous bias: the DM is biased towards a_1 because it is ex ante more likely under optimal inquiry, but it is ex ante more likely because the DM is biased towards it, even though this choice may be suboptimal ex post.

6. CASE STUDIES

We offer two case studies that illustrate the potential applications of our model. There are two important social institutions where information processing primarily takes the form of explicit inquiries: doctor visits and criminal investigations. In recent years, research has indicated that misdiagnosis and wrongful convictions, both

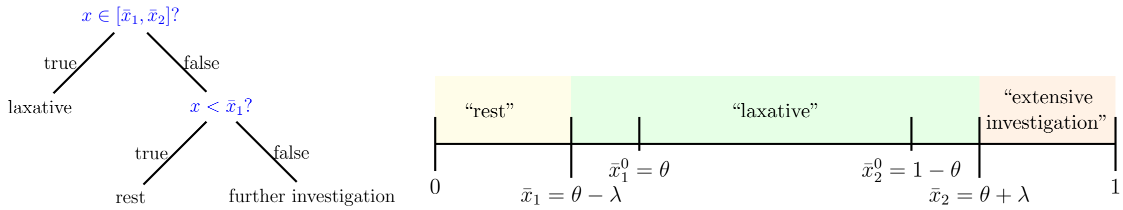

FIGURE 10. Optimal doctor inquiry tree (left) and optimal categories (right)

of which impact the life of the people affected significantly, are in fact closely linked to cognitive factors in the inquiries.

6.1. **Medical Misdiagnosis.** Singh et al. (2017) argue that "...diagnosis in primary care (i.e., first-contact, accessible, continued, comprehensive and coordinated care) represents a high-risk area for errors. PCPs typically face high patient volumes and make decisions amid uncertainty." They claim that the amount of errors is significant: "...a recent study estimated that about 5% of US adult patients experience diagnostic errors (defined as missed opportunities to make a correct or timely diagnosis) [...] every year." They point out that diagnostic reasoning is an important factor: "Several experimental studies have highlighted reasoning biases, in relation to both hypothesis generation and information interpretation in PCP's." Croskerry et al. (2013) discuss the mode of diagnostic reasoning (which they call "type 2"), which is the more typical one and is subject to "biases": "Our systematic errors are termed biases and there are many of them–biases over a hundred cognitive and approximately one dozen or so affective biases (ways in which our feelings influence our judgment)."

Croskerry et al. (2013) also describe a case study, in which a patient complained about constipation but actually was suffering from Cauda Equina Syndrome. This situation may be examined using our one-dimensional example with three actions: send the patient home to rest, prescribe a laxative, or refer for an extensive investigation. The value of each action depends on the true condition (the state x) of the patient. Suppose that the utility function is as in Figure 1 (interpreting a_1 as "rest", a_2 as "laxative", and a_3 as "extensive investigation"), the state x is uniformly distributed, and $\theta < 1/3$, so that "laxative" is the most likely correct action ex ante.

For a moderate value of λ, the optimal inquiry would assign the category associated with prescribing laxative with inquiry length of one. That is, the doctor should prioritize to enquire about whether or not it is best just to prescribe laxative. We depict the optimal inquiry tree and the optimal categories in Figure 10, where $\bar{x}_1$ and $\bar{x}_2$ are thresholds between the categories under optimal inquiry, and $\bar{x}_1^0$ and $\bar{x}_2^0$ are thresholds under zero-cost benchmark. Theorem 5.1 predicts that the doctor would prescribe laxative on a larger set of states (the interval $[\bar{x}_1, \bar{x}_2]$), as compared to the benchmark where the state of the patient can be discovered at zero cost (the interval $[\bar{x}_1^0, \bar{x}_2^0]$). Thus, in the intervals of states $(\bar{x}_1, \bar{x}_1^0)$ and $(\bar{x}_2^0, \bar{x}_2))$, the doctor makes an error. This is consistent with the argument in Croskerry et al. (2013) that the error is due to the following: "The principle biases for the physician who saw him in the clinic were framing, search satisficing and premature diagnostic closure." In our model, framing and search satisficing can be explained by the confirmation bias we identified. Indeed, as in our model, e.g., for $x \in (\bar{x}_2^0, \bar{x}_2)$, the doctor could have continued the investigation and potentially reached the conclusion that a more extensive investigation is needed, but stopped prematurely and prescribed a laxative.

An advantage of our model is that we can define "bias" rigorously. In this example, when the doctor prescribes "laxative" instead of "extensive investigation" for $x \in (\bar{x}_2^B, \bar{x}_2)$, we may call it an *"error"* and claim that the process is "biased". However, it is the *process* that is biased but not the decision per se. Moreover, relative to other models of imperfect information processing, our model is able to make prediction about the inquiry process, in this case, about the "premature diagnosis," and link that to the behavioural biases endogenously.

6.2. **Wrongful Conviction.** Gould and Leo (2010) review the literature on the extent and factors leading into wrongful convictions and believe that it is the *process* and factors affecting the process that are important. In their words, "...it is better to understand the sources of wrongful convictions not so much as dichotomous causes – a witness correctly or incorrectly identified the defendant and the identification directly led the jury to convict – but as contributing factors in a path analysis that might have been broken at some point before conviction." Among the leading factors the article identifies, we are interested in "tunnel vision", which is described in Gould and Leo (2010) as "the more law enforcement practitioners become convinced of a conclusion – in this case, a suspect's guilt – the less likely they are to consider alternative scenarios that conflict with this conclusion."

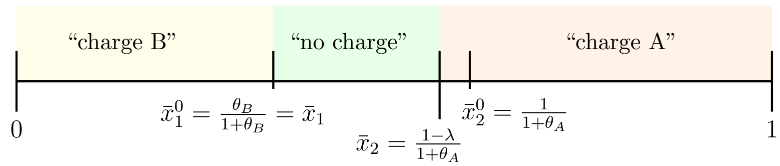

FIGURE 11. Optimal categories for the policy inquiry

We illustrate this tunnel vision with the following example. Suppose that there are two suspects, A and B, and one of them is surely guilty. Given all the possible observables the police can investigate, suppose that state $x \in [0, 1]$ represents the posterior belief that A is guilty. There are three actions: charge A, B, or neither, denoted by a_A, a_B, and $a_\varnothing$, respectively. Suppose that the police obtains utility θ_A if they charge A when A is guilty, θ_B if charge B when B is guilty, -1 if a wrong suspect (either A or B) is charged, and 0 if neither. Thus,

$$u(a_A, x) = \theta_A x + (-1)(1 - x), \quad u(a_B, x) = \theta_B(1 - x) + (-1)x, \quad u(a_\varnothing, x) = 0.$$

Assume that x is distributed according to the uniform distribution, and that $1 > \theta_A > \theta_B \geq 0.52$, so that ex ante the most likely suspect is A, and no action is dominated. We depict the optimal categories in Figure 11, where $\bar{x}_1$ and $\bar{x}_2$ are thresholds between the categories under optimal inquiry, and $\bar{x}_1^0$ and $\bar{x}_2^0$ are thresholds under zero-cost benchmark. As indicated in the figure, a positive λ leads to an expansion of the category for charging A, which was the primary suspect. In other words, the confirmation bias leads the police to lower the threshold of evidence needed to charge suspect A relative to the benchmark case without such cost. Specifically, on the interval of states $(\bar{x}_2, \bar{x}_2^0)$, the police makes an error by charging A when they should have let them go. Moreover, this interval expands with λ. This may be an explanation of the tunnel vision: the police under pressure to end the investigation optimally focuses on the prime suspect and is willing to charge the prime suspect even with relatively weak evidence.

Furthermore, if the cognitive cost were to rise even higher, the police would have found it optimal to drop the no-charge option out of the consideration set altogether. Thus, if we define "type-I" error as the situation where the police does not charge anyone, and "type-II" error as the situation where police charges the wrong suspect, then the optimal inquiry is always biased toward the type-II error. That is, it always has higher type-II error than the no-cost benchmark. This may be an explanation of the tunnel vision: the police under pressure to end the investigation optimally focuses on fewer options that what they would have considered with no such pressure.

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 (8). To simplify notation, let $X_n = I_n(Q^*)$ for each $n \in N$. By Lemma 3.1,

$$\mathbb{P}(X_n) > 0 \;\text{ for all } n \in N. \tag{19}$$

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

$$W_n(\hat{Q}; \lambda) = \frac{1}{\mathbb{P}(X_n)} \left(\sum_{t \in \hat{T}} \int_{x \in I_t(\hat{Q})} u(\hat{d}_t, x) - \lambda \ell_t(\hat{Q}) \right) G(\mathrm{d}x | X_n), \tag{20}$$

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 (20) 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 (20),

$$\sum_{t\in T_n} \left(\int_{x\in I_t(Q^*_n)} u(d_t, x) - \lambda \ell_t(Q^*_n) \right) G(\mathrm{d}x|X_n) \\ < \sum_{t\in\hat{T}} \left(\int_{x\in I_t(\hat{Q})} u(\hat{d}_t, x) - \lambda \ell_t(\hat{Q}) \right) G(\mathrm{d}x|X_n). \tag{21}$$

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 (19) and (20). 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*

$$\sum_{k=1}^{K} 2^{-\ell_k} = 1. \tag{22}$$

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

$$I_{t_k}(Q) = I_k \quad \text{and} \quad \ell_{t_k}(Q) = \ell_k \quad \text{for all } k = 1, ..., K. \tag{23}$$

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), (22) 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 (23). 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

$$I'_k = I_{t_k}(Q) \text{ for each } k = 1, ..., K-1, \quad \text{and} \quad I'_K = I_{t_K}(Q) \cup I_{t_{K+1}}(Q), \qquad (24)$$

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

$$\ell_{t_k}(Q) \leq \ell'_k \text{ for all } k = 1, ..., K, \ \ell_{t_k}(Q) < \ell'_k \text{ for some } k \in \{1, ..., K\}, \qquad (25)$$

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

$$I'_{t_k}(Q') = I'_k \text{ and } \ell'_{t_k}(Q') = \ell'_k \text{ for all } k = 1, ..., K. \tag{26}$$

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 (26), and the inequality is by (24), (25), 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 (8), for any I and any $d \in D$, if $x \in I^*_d(D, \ell) \cap I_{d'}$ with $d \neq d'$ 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. $\square$

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 (10).

A.7. **Proof of Proposition 3.1.** (a) If (D, ℓ) solves (10), 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. $\square$

A.8. **Proof of Theorem 4.1.** Let $\lambda_1 < \lambda_2$. For each $j = 1, 2$, Let Q_{λ_j} be an optimal inquiry for $j = 1, 2$, and let $Z_{\lambda_j} = (D^j, \ell^j, I^j)$ be the associated outcome. Denote

$$\bar{u}(Z_{\lambda_j}) = \sum_{d \in D^j} \int_{x \in I_d^j} u(d, x) G(\mathrm{d}x), \quad j = 1, 2.$$

By (7) and (12), we have $W(Q_{\lambda_j}; \lambda_j) = \bar{u}(Z_{\lambda_j}) - \lambda \bar{\ell}(Z_{\lambda_j})$. By the optimality of Z_{λ_j} given λ_j, for each $j = 1, 2$, we have

$$\bar{u}(Z_{\lambda_1}) - \lambda_1 \bar{\ell}(Z_{\lambda_1}) \geq \bar{u}(Z_{\lambda_2}) - \lambda_1 \bar{\ell}(Z_{\lambda_2}) \text{ and } \bar{u}(Z_{\lambda_2}) - \lambda_2 \bar{\ell}(Z_{\lambda_2}) \geq \bar{u}(Z_{\lambda_1}) - \lambda_2 \bar{\ell}(Z_{\lambda_1}).$$

Combining these inequalities yields

$$\lambda_1 \left(\bar{\ell}(Z_{\lambda_1}) - \bar{\ell}(Z_{\lambda_2})\right) \leq \bar{u}(Z_{\lambda_1}) - \bar{u}(Z_{\lambda_2}) \leq \lambda_2 \left(\bar{\ell}(Z_{\lambda_1}) - \bar{\ell}(Z_{\lambda_2})\right).$$

Thus, $\bar{\ell}(Z_{\lambda_1}) \geq \bar{\ell}(Z_{\lambda_2})$ whenever $\lambda_1 < \lambda_2$.

Next, let (D, ℓ) be a solution of problem (10) under $\lambda > 0$ such that ℓ is not uniform. Because (D, ℓ) is finite, there exists an interval $[\lambda', \lambda'']$ that contains λ

(possibly, $\lambda = \lambda'$ or $\lambda = \lambda''$) such that (D, ℓ) is a solution of problem (10) for each cost in $[\lambda', \lambda'']$. Consider arbitrary $\lambda_1, \lambda_2 \in [\lambda', \lambda'']$ with $\lambda_1 < \lambda_2$. Let $E_k^{\lambda_j}$ be given by (15). Observe that for each $j = 1, 2$ we have

$$\mathbb{P}(E_k^{\lambda_j}) = \mathbb{P}\left(\bigcup_{k'=1}^{k} I^*_{d_{k'}}(D, \ell; \lambda_j)\right) = \sum_{k'=1}^{k} \mathbb{P}(I^*_{d_{k'}}(D, \ell; \lambda_j)).$$

Thus, by Lemma A.1 (see Section A.11 below), for each $k = 1, ..., K-1$ we obtain

$$\sum_{k'=1}^{k} \mathbb{P}(I^*_{d_{k'}}(D, \ell; \lambda_2)) \geq \sum_{k=1}^{K} \mathbb{P}(I^*_{d_k}(D, \ell; \lambda_1)), \tag{27}$$

with strict inequality for some k. In other words, given (D, ℓ), the probability distribution over actions in $D = \{d_k\}_{k=1}^K$ under λ_2 first-order stochastically dominates that under λ_1. Thus, by (12), we obtain

$$\bar{\ell}(Z_{\lambda_2}) = \sum_{k=1}^{|D|} \ell_{d_k} \mathbb{P}(I^*_{d_k}(D, \ell; \lambda_2)) < \sum_{k=1}^{|D|} \ell_{d_k} \mathbb{P}(I^*_{d_k}(D, \ell; \lambda_1)) = \bar{\ell}(Z_{\lambda_1}).$$

Thus, we have shown that $\bar{\ell}(Z_\lambda)$ is strictly increasing on $[\lambda', \lambda'']$. □

A.9. **Proof of Proposition 4.1.** 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

$$I_{a'} = \{x \in X : u(a', x) > u(a, x) - \lambda(\ell_{a'} - \ell_a) \text{ for all } a \in D - \{a'\}\} \tag{28}$$

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 (28), 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

$$\ell_{a'} = \ell_{a''} = \hat{\ell}_{a'} + 1, \tag{29}$$

we obtain $2^{-\ell_{a'}} + 2^{-\ell_{a''}} = 2^{-\hat{\ell}_{a'}}$. 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 (7), the second equality is by (29), 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 4.2.** Let $\lambda_2 = \sup_{x\in X}(\max_{a\in A} u(a,x) - \min_{a\in A} u(a,x))$. Then, for all $\lambda \geq \lambda_2$, the utility gain from distinguishing any actions is smaller than the cost, so the optimal consideration set D is a singleton.

Next, by Theorem 3.2 and assumptions (A_1) and (A_3), there exists a nonempty set $\mathcal{F}_0 \subset \mathcal{F}^*$ of pairs (D,ℓ) that are optimal for a small enough cost. Specifically, for each $(D,\ell) \in \mathcal{F}_0$, there exists a small enough $\lambda_{(D,\ell)} > 0$ such that (D,ℓ) solves problem (10) for all $\lambda \in [0, \lambda_{(D,\ell)}]$. Let $\lambda_1 = \min_{(D,\ell)\in\mathcal{F}_0} \lambda_{(D,\ell)}$. Observe that by ($A_2$) and ($A_3$) and the assumption of full support on X, if $\lambda = 0$, then each action $a \in A$ is optimal on a positive-measure subset of X. Hence, the unique optimal consideration set for $\lambda = 0$ is $D = A$. It follows that $(D,\ell) \in \mathcal{F}_0$ implies $D = A$. Thus, A is the unique optimal consideration set for all $\lambda \in [0, \lambda_1]$. □

A.11. **Proof of Theorem 5.1.** Before proving Theorem 5.1, we state a lemma.

Lemma A.1. *Let $(D,\ell) \in \mathcal{F}^*$, and let $K = |D|$. W.l.o.g, let actions in D be ordered according to their lengths of inquiry, so $D = \{d_k\}_{k=1}^{K}$, such that*

$$\ell_{d_1} \leq \ell_{d_2} \leq \ldots \leq \ell_{d_{\bar{K}}}. \tag{30}$$

For each $\lambda_1, \lambda_2 \in \mathbb{R}_+$ *with* $\lambda_1 < \lambda_2$,

$$E_k^{\lambda_1} \subseteq E_k^{\lambda_2} \quad \textit{for all } k = 1, 2, ..., K-1. \tag{31}$$

Moreover, if ℓ *is not uniform, then there exists* $k \in \{1, ..., K-1\}$ *such that the set*

$$E_k^{\lambda_2} - E_k^{\lambda_1} \quad \textit{has a non-empty interior.} \tag{32}$$

Proof. First, we prove (31). Let $k \in \{1, ..., K-1\}$. Suppose by contradiction that there exists $x \in E_k^{\lambda_1}$ such that $x \notin E_k^{\lambda_2}$. By (15), $x \in E_k^{\lambda_1}$ and $x \notin E_k^{\lambda_2}$ imply that there exist $k^* \leq k < m^*$ such that

$$u(d_{k^*}, x) - \lambda_1 \ell_{d_{k^*}} = \max_{k'=1,...,k} u(d_{k'}, x) - \lambda_1 \ell_{d_{k'}} > u(d_{m^*}, x) - \lambda_1 \ell_{d_{m^*}}, \tag{33}$$

$$u(d_{k^*}, x) - \lambda_2 \ell_{d_{k^*}} \leq \max_{m=k+1,...,K} u(d_m, x) - \lambda_2 \ell_{d_m} = u(d_{m^*}, x) - \lambda_2 \ell_{d_{m^*}}. \tag{34}$$

Combining (33) and (34), we obtain

$$\lambda_2(\ell_{d_{m^*}} - \ell_{d_{k^*}}) \leq u(d_{m^*}, x) - u(d_{k^*}, x) < \lambda_1(\ell_{d_{m^*}} - \ell_{d_{k^*}}),$$

which is impossible, since $\lambda_2 > \lambda_1 \geq 0$, and $\ell_{d_{m^*}} \geq \ell_{d_{k^*}}$ by (30). We reached a contradiction.

Next, suppose that ℓ is not uniform. Then there exists $k^* \in \{1, ..., K-1\}$ such that

$$\ell_{d_1} = ... = \ell_{d_{k^*}} < \ell_{d_{k^*+1}} \leq ... \leq \ell_{d_K}. \tag{35}$$

We prove (32) for $k = k^*$. Define

$$\bar{w} = \max_{k=1,...,k^*} u(d_k, x) \quad \text{and} \quad w_\lambda(x) = \left(\max_{m=k^*+1,...,K} u(d_m, x) - \lambda \ell_{d_m} \right) - (\bar{w} - \lambda \ell_{d_{k^*}}).$$

Observe that

$$w_{\lambda_2}(x) < w_{\lambda_1}(x) \quad \text{for all } x \in X \text{ and all } \lambda_1 < \lambda_2. \tag{36}$$

This is because for any given $x \in X$ there exists $m^* > k^*$ such that

$$w_{\lambda_2}(x) = u(d_{m^*}, x) - \lambda_2 \ell_{d_{m^*}} + \lambda_2 \ell_{d_{k^*}} - \bar{w} < u(d_{m^*}, x) - \lambda_1 \ell_{d_{m^*}} + \lambda_1 \ell_{d_{k^*}} - \bar{w} \leq w_{\lambda_1}(x),$$

where the strict inequality is by $\lambda_1 < \lambda_2$ and $\ell_{k^*} < \ell_{m^*}$.

Next, by (15) and (35), we have

$$x \in E_{k^*}^{\lambda} \iff w_\lambda(x) < 0. \tag{37}$$

Fix $\lambda_1 < \lambda_2$. By assumptions (A_2)–(A_3), the sets $E_{k^*}^{\lambda_1}$ and $X - E_{k^*}^{\lambda_2}$ have nonempty interiors. Let

$$y \in Int(E_{k^*}^{\lambda_1}) \quad \text{and} \quad z \in Int(X - E_{k^*}^{\lambda_2}).$$

By (36) and (37), we have

$$w_{\lambda_1}(y) < 0 < w_{\lambda_2}(z) < w_{\lambda_1}(z).$$

Let

$$x^* = \alpha^* y + (1 - \alpha^*)z, \quad \text{where } \alpha^* = \sup\left\{\alpha \in [0,1] : w_{\lambda_1}(\alpha y + (1-\alpha)z) \leq 0\right\}.$$

Since X is convex, and points y and z are in $Int(X)$, x^* is an interior point of X. Since $w_{\lambda_1}(x)$ is continuous in x by assumption (A_1), we have $w_{\lambda_1}(x^*) = 0$. Moreover, by (36), there exists $\varepsilon > 0$ such that $w_{\lambda_2}(x^*) = \varepsilon > 0$. Let O_{x^*} be the open neighborhood of x^* given by

$$O_{x^*} = \{x \in X : |w_{\lambda_1}(x) - w_{\lambda_1}(x^*)| < \varepsilon\}.$$

By the continuity of $w_{\lambda_1}(x)$, O_{x^*} is an open nonempty set. Recall that by assumption (A_3), $(X - E_{k^*}^{\lambda_1})$ has nonempty interior. Since the set $O_{x^*} \cap (X - E_{k^*}^{\lambda_1})$ contains x^*, it has nonempty interior. Finally, since $O_{x^*} \cap (X - E_{k^*}^{\lambda_1}) \subset E_{k^*}^{\lambda_2} - E_{k^*}^{\lambda_1}$, we obtain (32) for $k = k^*$. □

We now prove Theorem 5.1. Let $\lambda > 0$ and let (D, ℓ, I) be an outcome of an optimal inquiry. Observe that $D = \{d_k\}_{k=1}^K$ satisfies (14) if and only if it satisfies (30). Then, the statement of the theorem is immediate by Definition 5.1 and Lemma A.1 with $\lambda_1 = 0$ and $\lambda_2 = \lambda$. □

A.12. **Proof of Proposition 5.1.** To prove the proposition, we apply Lemma A.1. Let $\lambda > 0$ and let (D, ℓ) be a solution of problem (10) under λ. Because (D, ℓ) is finite, there exists an interval $[\lambda', \lambda'']$ that contains λ (possibly, $\lambda = \lambda'$ or $\lambda = \lambda''$) such that (D, ℓ) is a solution of problem (10) for each cost in $[\lambda', \lambda'']$. Consider arbitrary $\lambda_1, \lambda_2 \in [\lambda', \lambda'']$ with $\lambda_1 < \lambda_2$. By Lemma A.1, it is immediate that, for each $k = 1, ..., K-1$,

$$\mathbb{P}(E_k^{\lambda_2}) \geq \mathbb{P}(E_k^{\lambda_1}). \tag{38}$$

Moreover, if ℓ is not uniform, inequality (38) is strict for some k. We thus obtain that confirmation bias is locally increasing, and strictly so whenever ℓ is not uniform. □

References

Bloedel, A. W. and W. Zhong (2021): "The cost of optimally-acquired information," *Working Paper.* 5

Caplin, A., M. Dean, and J. Leahy (2019): "Rational inattention, optimal consideration sets, and stochastic choice," *Review of Economic Studies*, 86, 1061–1094. 6, 10

Cover, T. M. and J. A. Thomas (2006): *Elements of Information Theory*, John Wiley & Sons, 2nd ed. 5

Cremer, J., L. Garicano, and A. Prat (2007): "Language and the Theory of the Firm," *The Quarterly Journal of Economics*, 122, 373–407. 5

Croskerry, P., G. Singhal, and S. Mamede (2013): "Cognitive debiasing 1: origins of bias and theory of debiasing," *BMJ Quality & Safety*, 22, 58–64. 10, 12, 33, 34

Denti, T., M. Marinacci, and A. Rustichini (2022): "Experimental cost of information," *American Economic Review*, 112, 3106–3123. 3

Gould, J. B. and R. A. Leo (2010): "One hundred years later: Wrongful convictions after a century of research," *The Journal of Criminal Law and Criminology*, 825–868. 34

Huffman, D. A. (1952): "A Method for the Construction of Minimum-Redundancy Codes," *Proceedings of the IRE*, 40, 1098–1101. 3

Jehiel, P. and J. Steiner (2020): "Selective sampling with information-storage constraints," *The Economic Journal*, 130, 1753–1781. 6

Jung, J., J. H. Kim, F. Matějka, and C. A. Sims (2019): "Discrete actions in information-constrained decision problems," *The Review of Economic Studies*, 86, 2643–2667. 5, 10

Kraft, L. G. (1949): "A device for quantizing, grouping, and coding amplitude-modulated pulses," Ph.D. thesis, Massachusetts Institute of Technology. 3

Maćkowiak, B., F. Matějka, and M. Wiederholt (2023): "Rational inattention: A review," *Journal of Economic Literature*, 61, 226–273. 2

Manzini, P. and M. Mariotti (2014): "Stochastic choice and consideration sets," *Econometrica*, 82, 1153–1176. 6

Masatlioglu, Y., D. Nakajima, and E. Y. Ozbay (2012): "Revealed attention," *American Economic Review*, 102, 2183–2205. 6

Matějka, F. and A. McKay (2015): "Rational inattention to discrete choices: A new foundation for the multinomial logit model," *American Economic Review*, 105, 272–298. 5, 10, 32

Morris, S. and P. Strack (2019): "The Wald problem and the equivalence of sequential sampling and static information costs," Available at SSRN: https://ssrn.com/abstract=2991567. 5

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

Nimark, K. P. and S. Sundaresan (2019): "Inattention and belief polarization," *Journal of Economic Theory*, 180, 203–228. 6

Pomatto, L., P. Strack, and O. Tamuz (2023): "The cost of information: The case of constant marginal costs," *American Economic Review*, 113, 1360–1393. 5

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

Singh, H., G. D. Schiff, M. L. Graber, I. Onakpoya, and M. J. Thompson (2017): "The global burden of diagnostic errors in primary care," *BMJ quality & safety*, 26, 484–494. 33

Steiner, J., C. Stewart, and F. Matějka (2017): "Rational inattention dynamics: Inertia and delay in decision-making," *Econometrica*, 85, 521–553. 6

Wald, A. (1945): "Sequential Tests of Statistical Hypotheses," *Annals of Mathematical Statistics*, 16, 117 – 186. 5

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