

NBER WORKING PAPER SERIES

PLANNING UNDER DIAGNOSTIC UNCERTAINTY:
QUESTION-DRIVEN LEARNING IN THE AGE OF AI

Andrew Caplin

Working Paper 35012
<http://www.nber.org/papers/w35012>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2026

I thank the Alfred P. Sloan Foundation for supporting this research under the Cognitive Economics at Work grant to NYU. Language models were used to support iterative inquiry and drafting; all research design, model construction, and conclusions are the author's. I am grateful to Søren Leth-Petersen and Ruth Wyatt for helpful discussions and comments. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Andrew Caplin. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Planning under Diagnostic Uncertainty: Question-Driven Learning in the Age of AI
Andrew Caplin
NBER Working Paper No. 35012
March 2026
JEL No. D81, D83, D91

ABSTRACT

Many important economic decisions depend not only on acquiring information, but on determining which questions are worth asking before committing to an action. Individuals and organizations must organize inquiry: they select questions, interpret answers, and update beliefs in ways that determine whether decisive distinctions are identified or missed. Despite its importance, the organization of inquiry is largely absent from formal economic analysis.

This paper studies decision-making under \emph{diagnostic uncertainty}, in which agents must learn not only about payoff-relevant states, but about which questions are informative. We develop a model in which agents choose questions sequentially prior to commitment, while facing uncertainty over a latent diagnostic structure that governs how questions generate answers. Before each inquiry step, agents may incur cognitive cost to refine beliefs about this diagnostic structure. The resulting problem is a finite-horizon dynamic program over inquiry, in which costly attention is allocated to belief transformations over diagnostic structure rather than directly to payoff-relevant states.

Two canonical diagnostic geometries isolate distinct planning margins. In one, value depends on locating a decisive question; in the other, on maintaining a correct sequence of questions. Both environments admit closed-form solutions and yield a common representation in which optimal inquiry increases the probability of success relative to uninformed search.

The framework identifies a distinct margin of economic behavior—planning under diagnostic uncertainty—that becomes increasingly important in environments where answers are abundant but the organization of inquiry remains scarce.

Andrew Caplin
New York University
Department of Economics
and NBER
andrew.caplin@nyu.edu

Planning under Diagnostic Uncertainty: Question-Driven Learning in the Age of AI

Andrew Caplin*

March 22, 2026

Abstract

Many important economic decisions depend not only on acquiring information, but on determining which questions are worth asking before committing to an action. Individuals and organizations must organize inquiry: they select questions, interpret answers, and update beliefs in ways that determine whether decisive distinctions are identified or missed. Despite its importance, the organization of inquiry is largely absent from formal economic analysis.

This paper studies decision-making under *diagnostic uncertainty*, in which agents must learn not only about payoff-relevant states, but about which questions are informative. We develop a model in which agents choose questions sequentially prior to commitment, while facing uncertainty over a latent diagnostic structure that governs how questions generate answers. Before each inquiry step, agents may incur cognitive cost to refine beliefs about this diagnostic structure. The resulting problem is a finite-horizon dynamic program over inquiry, in which costly attention is allocated to belief transformations over diagnostic structure rather than directly to payoff-relevant states.

Two canonical diagnostic geometries isolate distinct planning margins. In one, value depends on locating a decisive question; in the other, on maintaining a correct sequence of questions. Both environments admit closed-form solutions and yield a common representation in which optimal inquiry increases the probability of success relative to uninformed search.

The framework identifies a distinct margin of economic behavior—planning under diagnostic uncertainty—that becomes increasingly important in environments where answers are abundant but the organization of inquiry remains scarce.

1 Introduction

Many economic decisions require not only acquiring information, but determining how to generate information that is relevant for action. Before committing to a choice, individuals and organizations must decide where to look, which sources to engage, and how to interpret what they find. The

*I thank the Alfred P. Sloan Foundation for supporting this research under the Cognitive Economics at Work grant to NYU. Language models were used to support iterative inquiry and drafting; all research design, model construction, and conclusions are the author's. I am grateful to Søren Leth-Petersen and Ruth Wyatt for helpful discussions and comments.

central difficulty is often not only uncertainty about the world, but uncertainty about how to approach that uncertainty in a disciplined way.

A simple example illustrates the point. An individual considering a transition into data science may gather extensive information about salaries, locations, programming languages, and training opportunities. Yet this accumulation of answers may still fail to resolve the decision. What matters is not only how much information is obtained, but whether inquiry succeeds in identifying the constraint that determines what is feasible. The problem is therefore not only to learn about the world, but to determine how to learn what matters about the world.

This type of problem is common, but it sits awkwardly within standard economic frameworks. Search models, bandit models, and rational inattention models provide powerful accounts of learning under uncertainty, but they typically operate within a known informational structure (Mortensen and Pissarides, 1994; Gittins, 1979; Berry and Fristedt, 1985; Sims, 2003; Matějka and McKay, 2015; Caplin and Dean, 2015; Caplin et al., 2022). Agents may choose how intensively to search, which experiment to sample, or how much attention to allocate, yet the mapping from actions to information is already specified. The problem is therefore to learn within a given informational environment, not to determine how that environment should be interrogated in the first place. This paper focuses on that earlier margin.

The paper also builds on a broader line of work that treats cognitive processes—such as belief formation, attention allocation, and the generation of mistakes—as measurable economic objects, and that emphasizes the design of data and environments capable of making such processes observable (Caplin, 2025a,b, 2023). Within that perspective, the organization of inquiry is not merely an informal prelude to decision-making. It is itself an economically consequential process that can, in principle, be modeled and measured.

The present paper studies decision-making under *diagnostic uncertainty*, in which agents are uncertain not only about payoff-relevant states, but about how inquiry reveals those states. The central object is a latent diagnostic structure governing how acts of inquiry generate informative signals. Agents must therefore learn at two levels: about the world, and about the informational structure through which the world can be learned. The distinction from existing models is not that signals are rigid rather than flexible. It is that the informational structure itself is not taken as known.

We develop a model in which agents sequentially select inquiries prior to commitment while facing uncertainty over this latent diagnostic structure. At each stage, agents may incur cognitive cost to refine beliefs about how inquiry is likely to be informative, and then choose how to proceed. The resulting problem is a finite-horizon dynamic program over inquiry, in which attention is allocated to belief transformations over diagnostic structure rather than directly to payoff-relevant states.

To isolate the implications of diagnostic uncertainty, the paper analyzes two canonical diagnostic geometries. In the first, value depends on locating a decisive inquiry that resolves uncertainty. In the second, value depends on maintaining a correct sequence of inquiry, where early missteps reduce

the value of later investigation. Both environments admit closed-form solutions and yield a common representation in which optimal inquiry increases the probability of success relative to uninformed search.

The framework also provides a basis for measurement. Because inquiry is represented as a sequence of observable actions, it becomes possible to record how individuals generate information prior to decision, rather than inferring that process only from beliefs or realized outcomes. This creates a link between formal models of information acquisition and empirical data on inquiry behavior.

Recent technological developments make this margin increasingly important. Artificial intelligence systems can generate candidate answers to a wide range of questions at very low cost (OpenAI, 2023). As the cost of answers declines, the constraint on decision-making shifts toward the organization of inquiry. The ability to identify informative questions and sequence them effectively becomes more valuable, and differences in planning discipline may have larger economic consequences. At the same time, interactions with AI systems produce observable traces of inquiry, creating new opportunities to measure how individuals organize investigation in practice (Noy and Zhang, 2023; Brynjolfsson et al., 2023).

The remainder of the paper develops these ideas formally. Section 2 introduces the inquiry environment and defines questions, diagnostic structures, and belief dynamics. Section 3 introduces costly diagnostic refinement. Section 4 embeds these objects in a dynamic program describing sequential inquiry before commitment. Section 6 analyzes two canonical diagnostic geometries. Section 7 discusses measurement using inquiry traces. Section 8 examines implications for AI-mediated decision environments. Section 9 situates the framework within related literatures.

Contribution. This paper introduces a model of information acquisition under diagnostic uncertainty, in which agents are uncertain not only about payoff-relevant states but about how inquiry generates informative signals. In contrast to standard models that operate within a known informational structure, the framework allows for uncertainty over the diagnostic structure itself and models how agents allocate costly attention to refining it. The result is a tractable dynamic model of inquiry with closed-form solutions in canonical environments. The framework also provides a basis for measurement by representing inquiry as a sequence of observable actions, allowing the preparation of decisions to be studied directly.

2 The Architecture of Inquiry

Organizing inquiry. Many consequential decisions require information gathering before commitment. Classical models of information acquisition typically assume that the informational properties of available signals are known. Agents may choose which signals to observe, but they do not need to determine how those signals should be interpreted.

In many environments, however, the problem begins one step earlier. Decision makers face

uncertainty not only about the state of the world but also about which questions are informative and how answers should be interpreted. Some inquiries reveal decisive distinctions; others produce information that is descriptive but irrelevant to the choice at hand.

Diagnostic Uncertainty

An environment exhibits *diagnostic uncertainty* when the informational value of available questions is itself uncertain. The agent is uncertain not only about the payoff-relevant state, but also about which inquiries are informative and how their answers should be interpreted.

When diagnostic uncertainty is present, planning becomes the problem of organizing inquiry. Before acting, the agent must decide which questions to ask, how to interpret the answers they generate, and how to sequence inquiry prior to commitment.

The framework developed here has three layers. First, the inquiry environment specifies actions, payoff states, questions, answers, and diagnostic structures. Second, the agent may engage in costly diagnostic reasoning about which questions are likely to be informative. Third, inquiry unfolds sequentially in a dynamic program prior to commitment.

Primitives of the environment. The environment consists of

- a finite set of actions $\mathcal{Y}$,
- a payoff state $\theta \in \Theta$,
- a finite set of admissible questions $\mathcal{Q}$,
- answer sets $\{A_q\}_{q \in \mathcal{Q}}$,
- diagnostic structures $D \in \mathcal{D}$, and
- a finite inquiry horizon T .

The agent holds beliefs over the joint state (θ, D) and may pose questions sequentially before committing to an action.

The variable θ determines payoffs, while D governs how questions generate answers. The pair (θ, D) therefore summarizes both what is true and how inquiry reveals it.

Questions as atomic units of inquiry

The model restricts learning to a finite set of admissible *questions*. A question is an act of inquiry that jointly specifies both the distinction being investigated and the source from which an answer is obtained. Consulting one advisor rather than another, querying one dataset rather than another, or posing the same prompt to different language models may correspond to different questions.

This restriction is a modeling device. It does not require that decision makers explicitly enumerate questions, only that their learning process can be represented as selecting among a set of available inquiry acts, whether consciously or implicitly.

2.1 Conceptual Primitives of Inquiry

Three objects organize the model: questions, answers, and diagnostic structures.

The central modeling step is to treat questions as primitive acts of inquiry. Rather than choosing abstract signals, the agent selects from a set of available inquiry acts that differ in what they investigate and how answers are generated. This representation captures the fact that learning is constrained not only by what is unknown, but also by which forms of inquiry are available.

Definition 1 (Question). A *question* is an admissible act of inquiry

$$q \in \mathcal{Q}$$

that specifies how a distinction is investigated and from which source an answer is obtained, generating a realized answer

$$a \in A_q.$$

Formally, questions are characterized by answer-generating distributions

$$P(a \mid \theta, D, q).$$

Definition 2 (Answer). An *answer* is the realized outcome

$$a \in A_q$$

generated when question q is posed. Answers provide information about both the payoff state θ and the diagnostic structure D .

This interpretation is important because inquiry may take many forms, including internal reflection, consultation with advisors, or interaction with data and computational systems. Such processes may be conscious or subconscious. The model abstracts from these implementation details and treats questions as the primitive units through which inquiry proceeds.

Definition 3 (Diagnostic Structure). A *diagnostic structure* $D \in \mathcal{D}$ specifies how answers generated by questions relate to the payoff state. For each question $q \in \mathcal{Q}$, it determines the distribution

$$P(a \mid \theta, D, q), \quad a \in A_q.$$

Different diagnostic structures correspond to different models of how answers should be interpreted.

Definition 4 (Diagnostic Uncertainty). An environment exhibits *diagnostic uncertainty* when the agent faces uncertainty over diagnostic structures

$$D \in \mathcal{D}$$

that govern how answers generated by questions relate to the payoff-relevant state.

2.2 Actions and Commitment

The agent ultimately commits to an action

$$y \in \mathcal{Y},$$

where $\mathcal{Y}$ is a finite set of feasible actions.

Payoffs depend on an unknown payoff state

$$\theta \in \Theta$$

through an expected utility function

$$U : \mathcal{Y} \times \Theta \rightarrow \mathbb{R}.$$

Utility is realized only upon commitment to an action. Inquiry itself generates no intrinsic payoff. The planning problem is therefore static in consequences but dynamic in the evolution of beliefs prior to commitment.

2.3 Questions as Informational Experiments

Before committing to an action, the agent may acquire information through inquiry. In information-theoretic terms, each question corresponds to an informational experiment whose outcomes are answers. When the diagnostic structure is known, asking question q induces a well-defined statistical experiment mapping payoff states to distributions over answers.

Definition 5 (Diagnostic Experiment). Fix a diagnostic structure D . Each question $q \in \mathcal{Q}$ induces a statistical experiment

$$\Pi_q^D : \Theta \rightarrow \Delta(A_q), \quad \theta \mapsto P(\cdot \mid \theta, D, q).$$

We refer to Π_q^D as the *diagnostic experiment* associated with question q under diagnostic structure D .

Conditional on D , each question therefore corresponds to a well-defined Blackwell experiment. Under diagnostic uncertainty, however, a question does not correspond to a single experiment, but to a mapping whose informational content depends on D . Hence it represents a *lottery over experiments* indexed by D .

Diagnostic structures may also differ in the range of possible answers generated by the same question. Formally, the support of $P(a \mid \theta, D, q)$ may depend on D , so that distinct diagnostic structures may generate different answer sets for the same inquiry act.




Figure 1: Diagnostic structure determines how inquiry generates information. A chosen question q interacts with the latent environment (θ, D) to produce an answer a . Different diagnostic structures D correspond to different informational experiments, so the same question may produce different answer distributions.

2.4 Beliefs and Bayesian Updating

The full state of the environment is

$$(\theta, D) \in \Theta \times \mathcal{D}.$$

The agent begins inquiry with a prior belief

$$\mu_0 \in \Delta(\Theta \times \mathcal{D}).$$

After asking question q and observing answer a , beliefs update according to Bayes' rule.

Bayesian update operator. For prior belief μ , question q , and answer a , define

$$B_{(\mu,q,a)} \in \Delta(\Theta \times \mathcal{D})$$

by

$$B_{(\mu,q,a)}(\theta, D) = \frac{\mu(\theta, D)P(a \mid \theta, D, q)}{\sum_{\tilde{\theta}, \tilde{D}} \mu(\tilde{\theta}, \tilde{D})P(a \mid \tilde{\theta}, \tilde{D}, q)}.$$

Answers therefore provide information about both the payoff state θ and the diagnostic structure D .

2.5 Subjective Representations

The framework is entirely subjective in the tradition of Savage (1954). The sets of admissible questions $\mathcal{Q}$, answer sets A_q , and diagnostic structures $\mathcal{D}$ represent the agent's perceived space of inquiry possibilities and interpretive models.

Different agents may entertain different questions, consider different possible answers, or posit different diagnostic structures governing how answers should be interpreted.

2.6 Inquiry Horizon

Inquiry unfolds over a finite horizon of T periods. In each period the agent may pose a question and observe the resulting answer. Only after the inquiry phase concludes does the agent commit to an action. The horizon T captures limited time, attention, and cognitive resources available for inquiry. We focus on environments in which the number of admissible questions exceeds the available inquiry horizon, so that the agent cannot exhaust all possible inquiries and must select which questions to ask.

From inquiry to diagnostic planning. The primitives introduced above separate two elements of information acquisition. Questions determine the acts of inquiry available to the agent, while diagnostic structures determine how the answers generated by those questions should be interpreted.

When the diagnostic structure is known, each question corresponds to a well-defined experiment. When it is uncertain, the agent must reason about which questions are likely to be informative. The next section introduces a formal representation of that reasoning. Diagnostic refinement captures costly cognition directed not toward the payoff state itself, but toward the structure that governs which inquiries are worth pursuing.

3 Diagnostic Planning

The preceding section defined the inquiry environment. Agents face an uncertain payoff state, may ask questions before committing to an action, and receive answers whose informational content depends on an uncertain diagnostic structure.

Those primitives describe how inquiry generates information. They do not yet describe how the agent decides *which questions are worth asking*. That is the task of the present section.

When the diagnostic value of available questions is uncertain, decision-making requires reasoning about the informational structure of inquiry itself. Agents must evaluate which distinctions are likely to be useful, which inquiries are likely to be decisive, and how answers may redirect subsequent investigation.

3.1 Relation to Existing Models

Many classical models of information acquisition assume that the informational structure itself is known.

Search models draw observations from known distributions. Bandit models treat each arm as a known experiment. Standard rational-inattention models allow agents to design signals about the state of the world but assume that the mapping from states to signals is known.

In all of these frameworks, the object of learning is the payoff state.

In contrast, the present framework allows the informational value of questions themselves to be uncertain. Agents must therefore allocate attention not only to learning about the state of the world but also to learning how inquiry reveals that state.

In this sense, Question-Driven Learning introduces a higher-order problem: the object of attention is not only the state of the world, but the informational structure governing how questions generate answers.

Classical information models	Question-Driven Learning
Signal or experiment structure known	Diagnostic structure uncertain
Learn about payoff state θ	Learn which questions reveal θ
Costs apply to acquisition or processing	Costs apply to diagnostic refinement
Experiments are given	Experiments must be identified
Question choice is downstream	Question choice is itself the central object

Table 1: Shift in the object of attention.

Shift in the domain of costly cognition

Classical models attach informational or cognitive cost to learning about payoff-relevant states. The present framework instead attaches cost to learning about *diagnostic structure*: determining which questions are informative and how inquiry should be organized.

The key modeling move is therefore to shift the domain of costly cognition from payoff states to diagnostic structure. Rather than allocating effort to learning what is true, the agent allocates effort to determining which inquiries are worth pursuing.

3.2 Two Forms of Learning

Two distinct forms of learning arise in the environment.

Learning about the payoff state. When a question is asked and an answer is observed, beliefs about the payoff-relevant state θ update according to Bayes' rule. This is the familiar form of information acquisition studied in search models, bandits, and rational-inattention environments.

Learning about diagnostic structure. Before selecting the next question, the agent may attempt to refine beliefs about which questions are likely to be informative. This corresponds to learning about the diagnostic structure D governing how questions generate answers.

The distinctive feature of the present framework is that these two forms of learning are separated. Updating after observing answers is treated as routine. The key scarce margin concerns costly refinement of beliefs about diagnostic structure.

3.3 Diagnostic Refinement

Let

$$\nu_t(D) = \sum_{\theta} \mu_t(\theta, D)$$

denote the marginal belief over diagnostic structures.

Before selecting the next question, the agent may refine beliefs about D . We represent this refinement as choosing a Bayes-plausible distribution over posterior beliefs about diagnostic structure,

$$\Gamma_t \in \mathcal{B}(\nu_t),$$

where

$$\mathcal{B}(\nu) = \left\{ \Gamma \in \Delta(\Delta(\mathcal{D})) : \sum_{\hat{\nu}} \Gamma(\hat{\nu}) \hat{\nu} = \nu \right\}.$$

Drawing $\hat{\nu}_t \sim \Gamma_t$ produces the refined belief

$$\gamma_t(\theta, D) = \mu_t(\theta \mid D) \hat{\nu}_t(D).$$

Refinement alters beliefs about which questions are informative without changing beliefs about payoff states conditional on diagnostic structure.

3.4 Cost of Diagnostic Refinement

Refining beliefs about diagnostic structure requires cognitive effort. We represent this effort by a cost functional

$$\mathcal{C}(\nu_t, \Gamma_t) \geq 0,$$

where the feasible refinement $\Gamma_t \in \mathcal{B}(\nu_t)$ depends on the current belief, a dependence we suppress in the notation.

This cost functional is defined for all belief states. In particular, for any belief ν that may arise during inquiry, the agent may select a Bayes-plausible refinement, and the associated cognitive cost is defined on this entire domain.

The degenerate refinement that leaves beliefs unchanged incurs no cost:

$$\mathcal{C}(\nu_t, \Gamma_{\nu_t}^\circ) = 0.$$

Costs therefore apply to transformations of beliefs in the diagnostic-belief domain rather than to learning about payoff states.

3.5 Diagnostic Refinement and Rational Inattention

Diagnostic refinement can be interpreted as a rational-inattention problem over diagnostic structures. The agent allocates attention across competing hypotheses about which questions are informative.

When refinement costs are low, agents can tailor their beliefs about diagnostic structure closely to the informational features of the environment. When refinement costs are high, beliefs remain coarse and question choice becomes less responsive to the true diagnostic structure.

3.6 Planning as Anticipatory Reasoning

In many real decision problems, individuals attempt to anticipate how inquiry may unfold before selecting the next question.

An entrepreneur deciding whether to launch a product may consider how different advisors are likely to respond and how those responses would shape the next step of investigation. Managers evaluating a strategic investment may anticipate how further analysis or consultation would alter their understanding of the problem. Users interacting with artificial intelligence systems often experiment with alternative prompts in order to determine which ones produce informative answers.

In each of these settings, planning involves reasoning about which questions are worth asking.

3.7 Planning Skill

Planning skill arises from the efficient allocation of diagnostic attention. Agents with low refinement costs can focus attention on distinctions that matter for selecting the next question. Agents with higher refinement costs rely on coarser beliefs about diagnostic structure.

- **Shallow refinement:** failure to recognize which questions are likely to be informative.
- **Mis-sequencing:** pursuit of diagnostically weak questions while more informative sources are overlooked.
- **Wasted anticipation:** cognitive effort is spent on contingencies that do not affect the next inquiry step.
- **Premature stopping:** inquiry is abandoned before useful sources of information are identified.
- **Diffuse exploration:** attention is spread across too many sources rather than concentrated on the most informative ones.

Two agents facing the same environment and sharing the same prior may therefore pursue very different inquiry sequences depending on how effectively they refine beliefs about which questions are informative.

3.8 From Diagnostic Planning to Sequential Inquiry

The preceding discussion isolates the central elements of diagnostic planning. Agents face uncertainty about both the payoff state and the diagnostic structure governing how questions generate answers. After observing an answer, beliefs about the payoff state update in the usual Bayesian manner. Before selecting the next question, however, agents may also refine beliefs about diagnostic structure by allocating cognitive effort to evaluating which questions are informative.

At each stage, the agent therefore faces two linked choices: how much effort to devote to refining beliefs about diagnostic structure, and which question to ask next. The value of asking a particular question depends not only on the information it may reveal about the payoff state, but also on the possibility that its answer may reshape the structure of subsequent inquiry.

The next section formalizes this sequential planning problem as a dynamic program. Because the value of a question depends on both the payoff state and the diagnostic structure, it will also be convenient to represent inquiry policies as specifying, for each possible pair (θ, D) , how questions are selected. This representation will be introduced in the next section.

4 Dynamic Structure of Inquiry

The preceding sections introduced the primitives of the inquiry environment and the diagnostic refinement operator. We now embed these objects in a finite-horizon dynamic program describing how inquiry unfolds before commitment. The dynamic program formalizes a simple idea: before each inquiry step, the agent may devote cognitive effort to clarifying which questions are worth asking, then select a question, observe an answer, and continue or stop.

4.1 Timing and State

There is a finite inquiry horizon of $T + 1$ questioning opportunities. Time is indexed by the number of inquiry opportunities remaining,

$$t \in \{0, 1, \dots, T\}.$$

Here $t = 0$ denotes the final inquiry period: the agent may still refine diagnostic beliefs and ask one final question before committing to an action.

At the beginning of period t the agent holds a belief

$$\mu_t \in \Delta(\Theta \times \mathcal{D}),$$

which summarizes all information generated by past inquiry.

Let

$$\nu_t(D) = \sum_{\theta} \mu_t(\theta, D), \quad \mu_t^{\Theta}(\theta) = \sum_D \mu_t(\theta, D)$$

denote the marginal beliefs over diagnostic structures and payoff-relevant states, respectively. The agent also faces a set of admissible questions

$$\mathcal{Q}_t \subseteq \mathcal{Q}.$$

Within each period, it is useful to distinguish three belief objects. The belief μ_t is the inherited prior entering period t . Diagnostic refinement then produces an intermediate belief γ_t , which governs question choice within the period. After a question is asked and an answer is observed, Bayesian updating produces the posterior belief μ_{t-1} carried into the next period. Thus the within-period belief transition is

$$\mu_t \rightarrow \gamma_t \rightarrow \mu_{t-1}.$$

It is important to distinguish the *current* question choice from the *full future inquiry plan*. In period t , conditional on continuation, the agent directly chooses only the next question to be asked. Future question choices are not fixed in advance, but are determined recursively after later answers generate new posterior beliefs.

4.2 Final Inquiry Period

The dynamic program is most clearly understood by starting with the final period $t = 0$.

In this period the agent may still refine diagnostic beliefs and ask one final question.

Diagnostic refinement. The agent chooses a refinement

$$\Gamma_0 \in \mathcal{B}(\nu_0).$$

A draw

$$\hat{\nu}_0 \sim \Gamma_0$$

produces the refined belief

$$\gamma_0(\theta, D) = \mu_0(\theta \mid D) \hat{\nu}_0(D).$$

Question choice and answer generation. Conditional on γ_0 , the agent selects a question

$$q \in \mathcal{Q}_0.$$

At this stage of the analysis, question choice is represented directly as the selection of a single admissible question. An answer

$$a \in A_q$$

is then realized according to

$$a \sim P(\cdot \mid \theta, D, q).$$

Posterior beliefs are obtained using the Bayesian update operator

$$\mu_{-1} = B_{(\gamma_0, q, a)}.$$

Terminal commitment. After observing the final answer the agent commits to an action

$$y \in \mathcal{Y}.$$

The terminal payoff is therefore

$$\bar{V}(\mu) = \max_{y \in \mathcal{Y}} \sum_{\theta} \mu^{\Theta}(\theta) U(y, \theta),$$

where μ^{Θ} denotes the marginal belief over payoff states.

For notational convenience define

$$V_{-1}(\mu, \mathcal{Q}) = \bar{V}(\mu).$$

4.3 General Inquiry Periods

Consider now a period $t \geq 1$. The agent begins with prior belief μ_t and marginal diagnostic belief ν_t .

Diagnostic refinement. The agent selects a refinement

$$\Gamma_t \in \mathcal{B}(\nu_t),$$

and a draw

$$\hat{\nu}_t \sim \Gamma_t$$

produces the refined belief

$$\gamma_t(\theta, D) = \mu_t(\theta \mid D) \hat{\nu}_t(D).$$

Question choice. Conditional on γ_t , the agent chooses a question

$$q \in \mathcal{Q}_t.$$

In this posterior-based formulation, the choice object is the next admissible question itself; later sections introduce an equivalent state-dependent stochastic choice representation that is more convenient for establishing existence of an optimal strategy.

Answer generation and belief updating. An answer

$$a \in A_q$$

is realized according to

$$a \sim P(\cdot \mid \theta, D, q).$$

Posterior beliefs are obtained through the update operator

$$\mu_{t-1} = B_{(\gamma_t, q, a)}.$$

The posterior belief becomes the prior entering the next inquiry period.

4.4 Continuation Value

For refined belief γ_t and question $q \in \mathcal{Q}_t$ define

$$W_{t-1}(\gamma_t, q, \mathcal{Q}_t) = \sum_{a \in A_q} P(a \mid \gamma_t, q) V_{t-1}(B_{(\gamma_t, q, a)}, \mathcal{Q}_t \setminus \{q\}),$$

where

$$P(a \mid \gamma_t, q) = \sum_{\theta, D} \gamma_t(\theta, D) P(a \mid \theta, D, q).$$

Thus $W_{t-1}(\gamma_t, q, \mathcal{Q}_t)$ is the value of choosing *question* q *next*, given refined belief γ_t , when all later inquiry choices will themselves be made optimally through the continuation value V_{t-1} .

4.5 Bellman Recursion

At the beginning of period t the agent may either commit immediately or continue inquiry.

Immediate commitment yields value

$$\bar{V}(\mu_t) = \max_{y \in \mathcal{Y}} \sum_{\theta} \mu_t^{\Theta}(\theta) U(y, \theta).$$

Continuation yields

$$J_t(\mu_t, \mathcal{Q}_t) = \sup_{\Gamma_t \in \mathcal{B}(\nu_t)} \left[-\mathcal{C}(\nu_t, \Gamma_t) + \mathbb{E}_{\hat{\nu}_t \sim \Gamma_t} \left(\max_{q \in \mathcal{Q}_t} W_{t-1}(\gamma_t, q, \mathcal{Q}_t) \right) \right].$$

The value function therefore satisfies

$$V_t(\mu_t, \mathcal{Q}_t) = \max \{ \bar{V}(\mu_t), J_t(\mu_t, \mathcal{Q}_t) \}.$$

This posterior-based recursion will later be recast in a form that represents question-selection behavior as a state-dependent stochastic choice rule over the underlying pair (θ, D) .

4.6 Within-Period Architecture for $t \geq 1$

For periods $t \geq 1$, the evolution of beliefs within a period follows the sequence

$$\mu_t \rightarrow \gamma_t \rightarrow \mu_{t+1}.$$

Because answers depend on the underlying pair (θ, D) , any reduced-form description of question choice must ultimately specify how the distribution of questions varies across possible (θ, D) realizations.

4.7 Existence and representation

For analytical convenience, we re-express the posterior-based dynamic program using a state-dependent stochastic choice (SDSC) representation. This is an equivalent formulation of the same inquiry problem in which the choice of the next question is represented as a state-contingent rule over (θ, D) .

The SDSC representation is introduced to express the cost of diagnostic refinement in a tractable way. Rather than working directly with belief-based controls, it records how optimal question choice varies across the underlying state. The construction follows the logic of the static rational-inattention representation of Matějka and McKay (2015), but is adapted here to a dynamic environment with endogenous information acquisition.




Figure 2: Within-period inquiry dynamics for $t \geq 1$.

Because the model is Bayesian, past inquiry matters only through the current belief. Accordingly, the SDSC formulation is indexed by the current belief state $\mu_t \in \Delta(\Theta \times \mathcal{D})$ and the current feasible question set $\mathcal{Q}_t \subseteq \mathcal{Q}$.

As noted above, the value of a question depends on the underlying pair (θ, D) . Any reduced-form description of question choice must therefore specify, for each possible (θ, D) , how questions are selected. For each $(\mu_t, \mathcal{Q}_t)$, define the SDSC domain

$$\mathcal{R}(\mathcal{Q}_t) \equiv \prod_{(\theta, D) \in \Theta \times \mathcal{D}} \Delta(\mathcal{Q}_t).$$

An SDSC rule is an element $\rho_t \in \mathcal{R}(\mathcal{Q}_t)$, that is, a collection of conditional choice probabilities

$$\rho_t = \{\rho_t(q \mid \theta, D)\}_{(\theta, D) \in \Theta \times \mathcal{D}}.$$

For each possible realization of the underlying state (θ, D) , the rule $\rho_t(\cdot \mid \theta, D)$ specifies a probability distribution over admissible questions.

This representation should be interpreted as a reduced-form description of behavior rather than

a literal conditioning on observed states. The agent does not observe (θ, D) directly. Instead, the SDSC rule summarizes how question choice varies with the underlying state, as induced by the agent's beliefs and decision problem.

Induced cost on the SDSC domain. In the posterior formulation, the cost of diagnostic refinement is

$$\mathcal{C}(\nu_t, \Gamma_t),$$

where ν_t is the diagnostic marginal of μ_t and $\Gamma_t \in \mathcal{B}(\nu_t)$ is a Bayes-plausible refinement. Because beliefs evolve endogenously, the Bellman recursion implicitly applies this same cost technology at every belief state.

In the SDSC formulation, this dependence must be made explicit. The cost of implementing a given SDSC rule depends on the current belief, because the set of feasible refinements is determined by ν_t .

Fix $(\mu_t, \mathcal{Q}_t)$ and define

$$\nu_t(D) = \sum_{\theta} \mu_t(\theta, D).$$

Let $\Gamma \in \mathcal{B}(\nu_t)$ be a feasible refinement, and let

$$\hat{\nu} \sim \Gamma$$

denote the realized posterior diagnostic belief. This induces the refined joint belief

$$\gamma(\theta, D) = \mu_t(\theta \mid D) \hat{\nu}(D).$$

In the posterior formulation, question choice is determined after diagnostic refinement, as a function of the refined belief γ . A choice rule maps refined beliefs into distributions over questions:

$$\alpha(\cdot \mid \gamma) \in \Delta(\mathcal{Q}_t).$$

That is, once a particular refined belief γ is realized, the agent selects a question according to the distribution $\alpha(\cdot \mid \gamma)$. The randomness in question choice therefore arises from two sources: the stochastic refinement of beliefs and the (possibly stochastic) choice rule conditional on those beliefs.

Given (Γ, α) , the induced joint distribution over latent states and question choices is

$$\mathbb{P}(q, \theta, D) = \sum_{\hat{\nu}} \alpha(q \mid \gamma) \Gamma(\hat{\nu}) \gamma(\theta, D),$$

where $\gamma(\theta, D) = \mu_t(\theta \mid D) \hat{\nu}(D)$. This expression aggregates over all possible refined beliefs $\hat{\nu}$ that may be generated by the refinement Γ . For each realization $\hat{\nu}$, the refinement induces a corresponding belief γ , which determines question choice through α , and assigns probability $\gamma(\theta, D)$

to each underlying state. The resulting distribution describes how often each question is asked at each underlying state once refinement and choice are combined.

Definition 6 (Consistency with a refinement). An SDSC rule ρ_t is *consistent* with a refinement $\Gamma \in \mathcal{B}(\nu_t)$ if there exists a choice rule α such that, for all (θ, D) and q ,

$$\rho_t(q \mid \theta, D) \mu_t(\theta, D) = \sum_{\hat{\nu}} \alpha(q \mid \gamma) \Gamma(\hat{\nu}) \gamma(\theta, D),$$

where $\gamma(\theta, D) = \mu_t(\theta \mid D) \hat{\nu}(D)$.

Consistency requires that the SDSC rule ρ_t can be generated by some underlying process in which beliefs are first refined according to Γ , and questions are then chosen according to α . In this sense, ρ_t is a reduced-form description of behavior that can be *implemented* by a refinement-and-choice pair (Γ, α) . The induced SDSC cost at $(\mu_t, \mathcal{Q}_t)$ is defined by

$$C^S(\rho_t; \mu_t, \mathcal{Q}_t) = \inf \{ \mathcal{C}(\nu_t, \Gamma) : \rho_t \text{ is consistent with } \Gamma \}.$$

Assumption 1 (Regularity of the induced SDSC cost). *For every $(\mu_t, \mathcal{Q}_t)$, the induced SDSC cost*

$$C^S(\cdot; \mu_t, \mathcal{Q}_t) : \mathcal{R}(\mathcal{Q}_t) \rightarrow \mathbb{R}_+$$

is finite, the infimum in its definition is attained, and C^S is continuous with respect to the standard Euclidean metric on $\mathcal{R}(\mathcal{Q}_t)$, viewed as a product of simplices over (θ, D) .

A sufficient condition is that the primitive refinement cost $\mathcal{C}(\nu, \Gamma)$ be posterior separable, with a continuous and convex aggregator over posterior beliefs (Caplin et al., 2022). This class of costs, developed includes standard information-theoretic specifications such as Shannon entropy and ensures that the induced SDSC cost is well behaved.

Law of motion of beliefs. Given $(\mu_t, \mathcal{Q}_t)$ and an SDSC rule ρ_t , the joint distribution over latent states, questions, and answers is

$$\mathbb{P}_t(\theta, D, q, a) = \mu_t(\theta, D) \rho_t(q \mid \theta, D) P(a \mid \theta, D, q).$$

This law coincides with the distribution induced by optimal posterior-based choice once the SDSC rule is consistent with some refinement. It induces a distribution over next-period beliefs as follows. For each realized pair (q, a) , define the posterior

$$\mu_{t+1} = B_{(\mu_t, q, a)}.$$

The probability of observing (q, a) under ρ_t is

$$\mathbb{P}_t(q, a) = \sum_{\theta, D} \mu_t(\theta, D) \rho_t(q \mid \theta, D) P(a \mid \theta, D, q).$$

Thus an SDSC rule determines a well-defined stochastic law of motion for beliefs through the sequence

$$(\theta, D) \rightarrow q \rightarrow a \rightarrow \mu_{t-1}.$$

Dynamic program. Define the value of immediate commitment

$$\bar{V}(\mu_t) = \max_{y \in \mathcal{Y}} \sum_{\theta} \mu_t^{\Theta}(\theta) U(y, \theta).$$

For $t \geq 0$, the SDSC continuation value is

$$\begin{aligned} J_t^S(\mu_t, \mathcal{Q}_t) = & \max_{\rho_t \in \mathcal{R}(\mathcal{Q}_t)} \left\{ -C^S(\rho_t; \mu_t, \mathcal{Q}_t) \right. \\ & \left. + \sum_{q \in \mathcal{Q}_t} \sum_{a \in A_q} \mathbb{P}_t(q, a) V_{t-1}^S(B_{(\mu_t, q, a)}, \mathcal{Q}_t \setminus \{q\}) \right\}. \end{aligned}$$

The value function satisfies

$$V_t^S(\mu_t, \mathcal{Q}_t) = \max \{ \bar{V}(\mu_t), J_t^S(\mu_t, \mathcal{Q}_t) \}.$$

Induced transition law. For each state $(\mu_t, \mathcal{Q}_t)$ and SDSC rule $\rho_t \in \mathcal{R}(\mathcal{Q}_t)$, let

$$\Lambda_t(\cdot \mid \mu_t, \mathcal{Q}_t, \rho_t)$$

denote the induced distribution over triples (q, a, μ_{t-1}) generated by any refinement-choice pair (Γ, α) that is consistent with ρ_t at $(\mu_t, \mathcal{Q}_t)$. By Definition 6, the induced joint law of (θ, D, q) is pinned down by ρ_t and μ_t , while the next-period belief μ_{t-1} is obtained from the posterior update associated with the underlying refinement. In the equivalence result below, the value of the SDSC problem is shown not to depend on which consistent pair is selected.

Theorem 1 (Existence and equivalence). *Suppose that Θ , $\mathcal{D}$, $\mathcal{Q}$, and each answer set A_q are finite, and that Assumption 1 holds.*

Then for every horizon $t \in \{-1, 0, 1, \dots, T\}$ and every admissible state $(\mu_t, \mathcal{Q}_t)$:

1. *The SDSC Bellman problem admits an optimal SDSC rule. In particular, there exists $\rho_t^* \in \mathcal{R}(\mathcal{Q}_t)$ attaining*

$$J_t^S(\mu_t, \mathcal{Q}_t).$$

2. The posterior/refinement and SDSC representations are equivalent:

$$J_t(\mu_t, \mathcal{Q}_t) = J_t^S(\mu_t, \mathcal{Q}_t), \quad V_t(\mu_t, \mathcal{Q}_t) = V_t^S(\mu_t, \mathcal{Q}_t).$$

3. Consequently, an optimal policy exists for the posterior Bellman problem

$$V_t(\mu_t, \mathcal{Q}_t) = \max \{ \bar{V}(\mu_t), J_t(\mu_t, \mathcal{Q}_t) \}.$$

Proof. See Appendix B. □

4.8 Joint Refinement Extension

For expositional clarity, refinement in the baseline model operates only on the diagnostic marginal ν_t . One could instead allow refinement over the full joint belief

$$\mu_t \in \Delta(\Theta \times \mathcal{D}),$$

with costs defined over Bayes-consistent distributions on this larger space.

Under analogous regularity conditions, the Bellman recursion and the one-step refinement result continue to hold. Restricting refinement to the diagnostic marginal therefore sharpens interpretation without limiting the generality of the framework.

Listening as the dual margin. The present formulation focuses on the organization of questions. An equivalent extension would allow costly refinement over the interpretation of answers. In that formulation the agent allocates attention not to selecting which question to ask next, but to determining which distinctions within a realized answer are diagnostically relevant. Planning would then involve the joint organization of questioning and listening.

This dual perspective highlights a broader point. Inquiry skill therefore has two complementary components: the ability to generate decisive questions and the ability to extract decisive distinctions from the answers they produce.

5 Diagnosticity in Lived Inquiry

The model developed above describes environments in which the diagnostic value of questions is uncertain. The following cases illustrate how this structure arises in practical decision problems.

In each case an answer plays two roles: it updates beliefs about the payoff-relevant state and simultaneously informs beliefs about the diagnostic value of available questions. Planning therefore consists in sequencing inquiry under uncertainty about both outcomes and sources of information.

5.1 Case I: Pure State Uncertainty — Career Offer Evaluation

An individual receives an initial job offer. The offer reveals a salary and role description, but long-term fit and growth prospects remain uncertain. The individual must decide whether to accept immediately, negotiate, or seek additional offers.

Suppose the informational properties of available inquiries are well understood. A follow-up conversation with the manager reveals promotion structure. A conversation with current employees reveals work culture. A background check of the firm’s financials reveals stability. The diagnostic value of each inquiry is known: each has predictable implications for long-run satisfaction and earnings.

The individual asks one question and receives an answer. That answer shifts beliefs about long-term fit and therefore changes which remaining inquiry is most valuable. Ordering matters because posterior beliefs about the payoff-relevant state alter the expected contribution of each subsequent question.

Formal mapping. Here, $|\mathcal{D}| = 1$. The experiment mapping Π_q is known ex ante. The payoff-relevant state θ encodes long-run job quality. Each inquiry updates $\mu(\theta)$. Planning corresponds to classical experiment selection over atomic questions. Sequencing is driven by belief over θ alone.

5.2 Case II: Pure Diagnostic Uncertainty — Choosing an Advisor

An entrepreneur is deciding whether to launch a new product. Several advisors are available, each offering strategic guidance. The entrepreneur is unsure which advisor’s advice is reliable in this context.

The entrepreneur consults Advisor A and receives an ambiguous recommendation. The content shifts beliefs slightly about demand, but more importantly, it raises doubt about Advisor A’s diagnostic value. Is this advice grounded in relevant expertise, or is it generic and uninformative?

The next step is not merely to ask a different question about demand. It is to decide whether to continue consulting Advisor A, switch to Advisor B, commission independent research, or redesign the investigative approach entirely. Answers reorganize beliefs about which source of inquiry is trustworthy.

Formal mapping. Here, θ represents market demand and is low-dimensional. The primary uncertainty concerns D , which indexes the diagnostic value of available questions (or advisors). Each answer updates the joint belief $\mu(\theta, D)$, with particular effect on beliefs about D . Sequencing concerns evaluation of the experiment mapping $\{\Pi_q^D\}_{q \in \mathcal{Q}}$ rather than refinement of θ alone.

5.3 Case III: Hybrid Planning — Choosing Among Human and AI Sources

A firm is evaluating whether to pivot its product strategy. Management consults three sources:

- An experienced industry advisor.

- An internal data analysis team.
- Two different large language models with distinct training characteristics.

Initial responses differ. The industry advisor provides high-level strategic concerns. The internal team produces quantitative projections with significant caveats. One language model generates optimistic scenarios; another produces more conservative assessments.

Each response shifts beliefs about the underlying payoff-relevant state θ (true profitability under alternative strategies). At the same time, discrepancies across sources raise uncertainty about diagnostic structure. Is the advisor overconfident? Is the internal data incomplete? Does one model systematically overstate upside? Should one source be weighted more heavily or excluded entirely?

Subsequent inquiry is shaped not only by revised beliefs about profitability, but by revised beliefs about which source of learning is diagnostically reliable.

Formal mapping. Here, admissible questions $q \in \mathcal{Q}$ include consulting each advisor, querying each model, or analyzing each dataset. The payoff-relevant state θ concerns true strategic profitability. The diagnostic regime D governs how each source maps states into answers. Each answer updates the joint belief $\mu(\theta, D)$. Sequencing reflects re-ranking of questions as beliefs about both profitability and source reliability evolve.

5.4 Common Structure

Across these environments, answers play a dual role. They update beliefs about the payoff-relevant state θ and simultaneously inform beliefs about the diagnostic value of available questions.

Planning therefore involves learning not only about the state of the world but also about which sources of inquiry are reliable. This dual learning problem is precisely the structure captured by the dynamic model developed above.

5.5 Search and Learning Models as Degenerate Cases

The structure above strictly contains many familiar models of search and learning prior to commitment.

When the experiment mapping is known ($|\mathcal{D}| = 1$), planning reduces to classical sequential experimentation. In multi-armed bandits, each question corresponds to an arm with a known sampling technology. In canonical job search, each observed offer answers a question drawn from a known distribution, and optimal stopping summarizes sequencing. In elimination-by-aspects and characteristics-based search, each attribute test is a question whose diagnostic properties are fixed.

In all such models, alternatives are uncertain but their meaning is known. The informational structure within which learning occurs is fixed ex ante.

Question-Driven Learning relaxes precisely this restriction. When the diagnostic value of questions is itself uncertain, planning must account not only for state uncertainty but for the possibility

that certain lines of inquiry reveal little. In that environment, question order becomes a distinct economic margin.

Standard search and experimentation models therefore arise as special cases of a richer architecture in which diagnostic structure is known.

5.6 Why Planning Is Difficult

When diagnostic structure is uncertain, the informational value of a question cannot be evaluated solely from its immediate effect on beliefs about θ . The agent must reason about which questions are likely to be informative.

In the advisor example, the entrepreneur must assess whether further consultation will clarify the reliability of the advisor. In the CV example, the candidate must decide whether refining a prompt will produce meaningful improvement. In the hybrid case, management must judge which analytical sources deserve attention.

Planning skill therefore lies in how effectively agents refine their beliefs about the diagnostic value of available questions before selecting the next inquiry step.

Why diagnostic learning is cognitively costly. When the experiment mapping is known, the agent can evaluate a candidate question using its immediate expected informational effect on beliefs about θ . When diagnostic structure is uncertain, this shortcut breaks. To decide what to ask next, the agent must assess how different answers would shift beliefs about the reliability of competing sources and thereby reorganize which questions are worth asking. This requires reasoning beyond the immediate posterior over θ and is naturally modeled as costly refinement of beliefs about D . We model this anticipatory diagnostic reasoning—not routine Bayesian updating—as the scarce cognitive input in planning.

The general model permits arbitrary diagnostic structures.

The environments described above illustrate the central challenge of diagnostic planning: determining which questions are informative when the diagnostic structure itself is uncertain.

From lived inquiry to economic geometry. The examples above illustrate a common planning difficulty. Decision makers must determine which questions are likely to generate useful information when the diagnostic structure of inquiry is uncertain.

The general model allows arbitrary diagnostic mappings

$$D \mapsto \{\Pi_q^D\}_{q \in \mathcal{Q}}.$$

In such environments inquiry can take many forms.

To clarify the economic forces governing optimal planning, the next section studies two simplified environments that isolate distinct planning margins.

The first geometry captures environments in which the key difficulty is identifying which question is decisive. The second captures environments in which value arises from maintaining the

correct sequence of inquiry.

These geometries abstract from the richness of the lived environments above, but they isolate the economic mechanisms that determine how diagnostic refinement affects optimal question choice.

6 Two Canonical Diagnostic Geometries

The dynamic model developed in the previous section allows the agent to refine beliefs about diagnostic structure before selecting a question. That formulation is intentionally general: the answer to a question may reshape the continuation problem in many different ways, and therefore different environments can generate very different patterns of inquiry.

To make the economic logic transparent, this section studies two canonical diagnostic geometries. In both cases the agent allocates attention to increase the probability that the next question successfully advances the diagnostic process. Preferences, timing, and the information technology are held fixed across environments. What differs across environments is how failure affects continuation value and therefore how inquiry should be organized.

Information acquisition follows the Shannon (mutual-information) specification. In the present framework, refinement operates on the diagnostic marginal belief

$$\nu_t(D) = \sum_{\theta} \mu_t(\theta, D).$$

Let $\Gamma_t \in \mathcal{B}(\nu_t)$ denote a Bayes-plausible distribution over posterior diagnostic beliefs $\hat{\nu}_t \in \Delta(\mathcal{D})$. Refining ν_t into Γ_t incurs cost

$$\kappa I(\Gamma_t; \nu_t) = \kappa \mathbb{E}_{\hat{\nu}_t \sim \Gamma_t} \left[\sum_{D \in \mathcal{D}} \hat{\nu}_t(D) \log \left(\frac{\hat{\nu}_t(D)}{\nu_t(D)} \right) \right].$$

The summation runs over diagnostic structures $D \in \mathcal{D}$. Thus the Shannon cost equals the expected Kullback–Leibler divergence between the refined diagnostic belief $\hat{\nu}_t$ and the prior diagnostic belief ν_t . In this formulation, belief refinement operates on the diagnostic marginal belief rather than the full joint belief over payoff and diagnostic states.¹

6.1 Economic environment and timing

The questions in the model matter only because they help the agent make a payoff-relevant decision. To make this explicit, suppose the agent ultimately faces a binary choice problem.

There are two equally likely payoff-relevant states,

$$\omega \in \{\omega_0, \omega_1\}, \quad \Pr(\omega_0) = \Pr(\omega_1) = \frac{1}{2}.$$

¹More generally, refinement could be defined over the full joint belief $\mu_t \in \Delta(\Theta \times \mathcal{D})$, yielding the standard Shannon cost $\mathbb{E}_{\hat{\mu} \sim \Gamma} [\sum_{\theta, D} \hat{\mu}(\theta, D) \log(\hat{\mu}(\theta, D)/\mu(\theta, D))]$. Restricting refinement to the diagnostic marginal simplifies the interpretation without affecting the structure of the model.

The agent must eventually choose an action

$$a \in \{a_0, a_1\},$$

where action a_j is correct in state ω_j . Utility is normalized so that a correct choice yields payoff 1 and an incorrect choice yields payoff 0:

$$u(a, \omega) = \mathbf{1}\{a \text{ matches } \omega\}.$$

The agent does not observe the state directly. Instead the agent has access to a collection of questions whose answers help organize the diagnostic process. Asking a question does not necessarily reveal the state itself. Rather, the answer may indicate which line of inquiry is productive or which candidate question should be asked next. The economic value of inquiry therefore lies in improving the probability that the eventual action matches the true state.

Periods are indexed by the number of remaining opportunities to refine beliefs and ask a question. If T periods remain, the agent has $T + 1$ opportunities to learn and choose a question, indexed

$$T, T - 1, \dots, 0.$$

Period 0 is the final learning stage before the last question is chosen.

At the beginning of period s the agent holds a belief over the currently relevant diagnostic structure. The agent may then choose an information structure—that is, a Bayes-consistent experiment over posterior beliefs—and pay the associated Shannon information cost. After observing the posterior belief the agent selects a question. The answer determines whether the prize is obtained immediately or whether the diagnostic process continues into the next period.

Learning is therefore available before action in every period, including the final one. What differs across diagnostic environments is how the answer to a question reshapes the continuation problem.

In the *identification geometry*, exactly one question resolves the payoff-relevant uncertainty. Asking the correct question immediately reveals the information needed for the final decision. Asking an incorrect question eliminates one candidate and the agent continues searching among the remaining possibilities.

In the *sequential-discipline geometry*, value arises from maintaining a correct sequence of questions. There exists a unique effective path of inquiry. Asking the correct next question preserves this path and allows the diagnostic process to continue, while asking any other question breaks the path and eliminates all future value.

These two environments differ only in how failure propagates: in the identification geometry, failure eliminates one candidate but preserves future value; in the sequential geometry, failure eliminates all future value.

The analysis in each geometry proceeds in the same way. We first describe the economic envi-

ronment and write the Bellman equation for the full inquiry problem allowing arbitrary refinement over diagnostic structure. We then introduce a restricted static problem in which the agent chooses directly the probability that the next question succeeds. This restricted problem admits a closed-form solution. The appendix then verifies that the candidate value function generated by the restricted problem satisfies the Bellman equation for the corresponding diagnostic geometry.

6.2 Identification Geometry

We begin with the identification geometry, in which inquiry proceeds by locating a single decisive question.

In the general model the state entering period s is a joint belief

$$\mu_s \in \Delta(\Theta \times \mathcal{D}),$$

with induced diagnostic marginal

$$\nu_s(D) = \sum_{\theta} \mu_s(\theta, D).$$

Let $n_s := |\mathcal{Q}_s|$ denote the number of feasible questions at stage s .

The relevant feature of the diagnostic state is the identity of the currently decisive question. Formally, for each diagnostic state D and current feasible question set $\mathcal{Q}_s$, let

$$d_s(D) \in \mathcal{Q}_s$$

denote the unique decisive question at stage s .

If the agent asks $d_s(D)$, the information needed for the final decision is obtained and the prize is realized. If the agent asks any other question, the answer is a failure signal: the chosen question is eliminated from consideration, and the diagnostic problem continues with the remaining candidates.

Thus the economic logic of the identification geometry is one of gradual elimination. Failure is informative but not fatal: it reduces the set of candidate questions while preserving the possibility of future success.

The full joint belief μ_s induces a marginal belief over which question is decisive. For each $q \in \mathcal{Q}_s$, define

$$\nu_s^I(q) = \sum_{\theta, D} \mu_s(\theta, D) \mathbf{1}\{d_s(D) = q\}.$$

This induced marginal summarizes the aspect of diagnostic uncertainty that matters for inquiry in the identification geometry.

At the beginning of period s , the agent may refine beliefs over the identity of the decisive question by choosing a Bayes-consistent experiment

$$\Gamma_s \in \mathcal{B}(\nu_s^I),$$

which yields a refined diagnostic belief

$$\hat{\nu}_s \sim \Gamma_s.$$

Conditional on $\hat{\nu}_s$, the agent chooses a question

$$q \in \mathcal{Q}_s.$$

If $q = d_s(D)$, which occurs with probability $\hat{\nu}_s(q)$, the prize is obtained immediately. If $q \neq d_s(D)$, which occurs with probability $1 - \hat{\nu}_s(q)$, the question fails, is eliminated, and the problem continues with feasible set

$$\mathcal{Q}_s \setminus \{q\}.$$

Conditional on failure of q , the posterior belief over the remaining decisive questions is the renormalized belief

$$\hat{\nu}_s^{-q}(q') = \frac{\hat{\nu}_s(q')}{1 - \hat{\nu}_s(q)}, \quad q' \in \mathcal{Q}_s \setminus \{q\}.$$

Under Shannon information costs, the specialized Bellman equation for the identification geometry is therefore

$$V_s^I(\nu_s^I, \mathcal{Q}_s) = \sup_{\Gamma_s \in \mathcal{B}(\nu_s^I)} \left[-\kappa I(\Gamma_s; \nu_s^I) + \mathbb{E}_{\hat{\nu}_s \sim \Gamma_s} \left(\max_{q \in \mathcal{Q}_s} \left\{ \hat{\nu}_s(q) + (1 - \hat{\nu}_s(q)) V_{s-1}^I(\hat{\nu}_s^{-q}, \mathcal{Q}_s \setminus \{q\}) \right\} \right) \right].$$

This is the Bellman equation for the identification geometry.

The analysis below focuses on the canonical symmetric benchmark in which all currently feasible questions are ex ante equiprobable under the induced diagnostic marginal. In that environment, and under the optimal refinement characterized in the appendix, the unrestricted Bellman problem collapses to a one-dimensional choice over the probability that the next question is decisive. This is the restricted sequence problem analyzed next.

6.2.1 Identification Geometry: A Restricted Sequence Problem

The next step is to represent this environment by a reduced problem in which the agent chooses directly the probability that the next question succeeds. We now introduce a restricted optimization problem that will serve as a tractable representation of the identification environment. This is a simple model in which the agent chooses directly the probability that the next question is decisive at each stage of the diagnostic process. The resulting problem is a classical finite-dimensional optimization problem whose solution can be characterized explicitly. We will show that the value function generated by this restricted problem also satisfies the Bellman equation of the full inquiry model.

Consider a diagnostic process with $T+1$ questioning opportunities. At stage s the agent chooses a conditional success probability

$$\hat{p}_s \in \left[\frac{1}{n_s}, 1 \right],$$

where n_s denotes the number of feasible candidate questions at that stage. The lower bound reflects the probability of success under uninformed search among the n_s candidate questions, while the upper bound corresponds to perfectly identifying the decisive question.

The sequence

$$(\hat{p}_T, \hat{p}_{T-1}, \dots, \hat{p}_0)$$

therefore summarizes a complete candidate strategy for the diagnostic process.

Success probability. In the identification geometry, the prize is obtained as soon as the decisive question q^* is asked. Hence the probability that the diagnostic process ultimately succeeds is

$$\Pr(\text{prize}) = 1 - \prod_{s=0}^T (1 - \hat{p}_s).$$

Define

$$\alpha_s = \prod_{r=s+1}^T (1 - \hat{p}_r),$$

the probability that stage s is reached—that is, the probability that all later attempts have failed.

Information costs. If conditional success probability $\hat{p}_s$ is implemented at a stage with n_s feasible candidate questions, the associated Shannon information cost is

$$C_s(\hat{p}_s) = \kappa \left[\hat{p}_s \log(n_s \hat{p}_s) + (1 - \hat{p}_s) \log\left(\frac{n_s(1 - \hat{p}_s)}{n_s - 1}\right) \right].$$

Because stage s is reached only with probability α_s , the expected information cost incurred at that stage is

$$\alpha_s C_s(\hat{p}_s).$$

The expected total information cost of the sequence is therefore

$$\sum_{s=0}^T \alpha_s C_s(\hat{p}_s).$$

The restricted optimization problem. The agent therefore chooses the sequence of conditional success probabilities to solve

$$\max_{(\hat{p}_T, \dots, \hat{p}_0)} \left\{ 1 - \prod_{s=0}^T (1 - \hat{p}_s) - \sum_{s=0}^T \alpha_s C_s(\hat{p}_s) \right\}.$$

This optimization problem determines the optimal organization of inquiry when the agent can directly control the probability that the next question is decisive at each stage. In the next subsection we solve this problem explicitly.

6.2.2 Identification Geometry: Closed-Form Solution

We now solve the restricted sequence problem introduced above.

Recall that the agent chooses the entire sequence of conditional success probabilities

$$(\hat{p}_T, \hat{p}_{T-1}, \dots, \hat{p}_0)$$

to maximize the objective defined in the previous subsection. This is a finite-dimensional optimization problem. The choice variables are the conditional success probabilities at each stage, and the objective combines the probability of eventually asking the decisive question with the expected Shannon information costs incurred along the way.

To compute the solution it is convenient to exploit the dependent structure of the objective function. In particular, the contribution of stage s depends on the probability that all later stages have failed. This allows the problem to be decomposed stage by stage and solved by backward induction.

Importantly, this recursion is simply a computational device for solving the restricted sequence optimization problem. The dynamic appearance arises only from the multiplicative structure of failure probabilities in the objective.

Let $\hat{U}_s$ denote the value of the restricted sequence problem when s questioning opportunities remain. Since no further success is possible after the last opportunity has passed, the terminal condition is

$$\hat{U}_{-1} = 0.$$

At a stage with s periods remaining, the agent chooses the conditional success probability $p \in [1/n_s, 1]$. If the next question is decisive, the prize is obtained immediately and the payoff is 1. If the question fails, the diagnostic process continues to the next stage with value $\hat{U}_{s-1}$. The stage- s problem is therefore

$$\hat{U}_s = \max_{p \in [1/n_s, 1]} \left\{ p + (1 - p)\hat{U}_{s-1} - \kappa \left[p \log(n_s p) + (1 - p) \log\left(\frac{n_s(1 - p)}{n_s - 1}\right) \right] \right\}.$$

The objective is strictly concave in p , so the optimum is unique. Differentiating with respect to p yields the first-order condition

$$1 - \widehat{U}_{s-1} = \kappa \log \left(\frac{p(n_s - 1)}{1 - p} \right).$$

Solving for the optimal conditional success probability gives

$$\hat{p}_s = \frac{1}{1 + (n_s - 1)e^{-(1 - \widehat{U}_{s-1})/\kappa}} = \frac{e^{1/\kappa}}{e^{1/\kappa} + (n_s - 1)e^{\widehat{U}_{s-1}/\kappa}}.$$

Substituting $\hat{p}_s$ back into the objective yields the recursion

$$\widehat{U}_s = \kappa \log \left(\frac{e^{1/\kappa} + (n_s - 1)e^{\widehat{U}_{s-1}/\kappa}}{n_s} \right).$$

This recursion can be solved explicitly. Define

$$U_s = e^{\widehat{U}_s/\kappa}.$$

Then

$$U_s = \frac{e^{1/\kappa} + (n_s - 1)U_{s-1}}{n_s}, \quad U_{-1} = 1.$$

Writing

$$B_s = \prod_{r=0}^s \frac{n_r - 1}{n_r},$$

the recursion solves to

$$U_s = e^{1/\kappa} - (e^{1/\kappa} - 1)B_s.$$

Hence the value of the restricted sequence problem is

$$\widehat{U}_s = \kappa \log \left(e^{1/\kappa} - (e^{1/\kappa} - 1) \prod_{r=0}^s \frac{n_r - 1}{n_r} \right).$$

In particular, with $T + 1$ questioning opportunities, the value of the restricted sequence problem is

$$\widehat{U}_T = \kappa \log \left(e^{1/\kappa} - (e^{1/\kappa} - 1) \prod_{r=0}^T \frac{n_r - 1}{n_r} \right).$$

This explicit expression is the value generated by the restricted sequence problem and will serve as the candidate value function in the verification theorem below.

6.2.3 Why the identification geometry remains tractable

The identification geometry is genuinely sequential. At each stage, the agent chooses a question, and the process stops as soon as the decisive question is found. What makes the restricted sequence problem tractable is the way Shannon information costs interact with this stopping structure.

If the question asked at stage s is decisive, the prize is obtained immediately. If it fails, the process continues to stage $s - 1$ with one fewer candidate question. Thus the probability of continuation multiplies not only future payoffs, but also all future optimized information costs. This is what allows the restricted problem to retain a recursive form.

In particular, the stage- s value can be written as

$$\hat{U}_s = \max_{p \in [1/n_s, 1]} \left\{ p + (1 - p)\hat{U}_{s-1} - \kappa \left[p \log(n_s p) + (1 - p) \log\left(\frac{n_s(1 - p)}{n_s - 1}\right) \right] \right\}.$$

The factor $(1 - p)$ is the probability that the decisive question has not yet been found. Because this same factor multiplies the entire continuation value $\hat{U}_{s-1}$, the optimized future payoff and future Shannon costs enter recursively as a single object. This recursive structure is what ultimately yields the closed-form solution

$$\hat{U}_T = \kappa \log \left(e^{1/\kappa} - (e^{1/\kappa} - 1) \prod_{r=0}^T \frac{n_r - 1}{n_r} \right).$$

Appendix 1 provides the short inductive calculation establishing this expression.

6.2.4 Verification theorem

The restricted sequence problem solved above yields the explicit value function

$$\hat{U}_T = \kappa \log \left(e^{1/\kappa} - (e^{1/\kappa} - 1) \prod_{r=0}^T \frac{n_r - 1}{n_r} \right).$$

We now verify that the candidate value function generated by the restricted sequence problem satisfies the Bellman equation of the full dynamic inquiry problem in the canonical symmetric environment.

Theorem 2 (Verification for the identification geometry). *Consider the canonical symmetric identification geometry in which, at each stage s , the induced diagnostic belief over the feasible questions $\mathcal{Q}_s$ is uniform.*

Let $\hat{U}_s$ denote the candidate value function generated by the restricted sequence problem. Then, for every stage s , the function $\hat{U}_s$ satisfies the unrestricted Bellman recursion for the identification geometry:

$$V_s^I(\bar{\nu}_s^I, \mathcal{Q}_s) = \hat{U}_s.$$

Moreover, at each stage the supremum in the unrestricted Bellman problem is attained by the

symmetric one-spike posterior family that implements the restricted success probability choice.

Hence the restricted sequence problem delivers the value of the full dynamic inquiry problem in the canonical symmetric identification environment.

The proof is given in Appendix C.1. It verifies directly that the candidate value function $\hat{U}_T$ is preserved by the Bellman operator.

6.3 Sequential-Discipline Geometry

We now consider the second canonical diagnostic environment. In this setting value arises from maintaining a correct sequence of questions rather than from locating a single decisive one.

In the general model the state entering period s is again a joint belief

$$\mu_s \in \Delta(\Theta \times \mathcal{D}),$$

with induced diagnostic marginal

$$\nu_s(D) = \sum_{\theta} \mu_s(\theta, D).$$

Recall $n_s := |\mathcal{Q}_s|$ denotes the number of feasible questions at stage s .

In the sequential geometry, the relevant feature of the diagnostic state is the identity of the unique question that preserves the effective path at the current stage. For each diagnostic state D and feasible question set $\mathcal{Q}_s$, let

$$c_s(D) \in \mathcal{Q}_s$$

denote the correct next question on the path.

If the agent asks $c_s(D)$, the path is preserved and inquiry continues to stage $s - 1$. If the agent asks any other question, the path is broken and all continuation value is lost.

The full joint belief μ_s therefore induces a marginal belief over which question is currently correct. For each $q \in \mathcal{Q}_s$, define

$$\nu_s^S(q) = \sum_{\theta, D} \mu_s(\theta, D) \mathbf{1}\{c_s(D) = q\}.$$

This induced marginal summarizes the aspect of diagnostic uncertainty that matters for inquiry in the sequential geometry.

At the beginning of period s , the agent may refine beliefs over the identity of the correct next question by choosing a Bayes-consistent experiment

$$\Gamma_s \in \mathcal{B}(\nu_s^S),$$

which yields a refined diagnostic belief

$$\hat{\nu}_s \sim \Gamma_s.$$

Conditional on $\hat{\nu}_s$, the agent chooses a question

$$q \in \mathcal{Q}_s.$$

If $q = c_s(D)$, which occurs with probability $\hat{\nu}_s(q)$, the path is preserved and the problem continues. If $q \neq c_s(D)$, which occurs with probability $1 - \hat{\nu}_s(q)$, the path is broken and continuation value falls to zero.

Let

$$S_s(\hat{\nu}_s, q)$$

denote the induced next-stage diagnostic belief conditional on successfully choosing question q . Under Shannon information costs, the specialized Bellman equation for the sequential geometry is therefore

$$V_s^S(\nu_s^S, \mathcal{Q}_s) = \sup_{\Gamma_s \in \mathcal{B}(\nu_s^S)} \left[-\kappa I(\Gamma_s; \nu_s^S) + \mathbb{E}_{\hat{\nu}_s \sim \Gamma_s} \left(\max_{q \in \mathcal{Q}_s} \left\{ \hat{\nu}_s(q) V_{s-1}^S(S_s(\hat{\nu}_s, q), \mathcal{Q}_{s-1}) \right\} \right) \right].$$

This is the unrestricted Bellman problem for the sequential geometry.

The analysis below focuses on the canonical symmetric benchmark in which all currently feasible questions are ex ante equiprobable and, conditional on choosing the correct question, the next-stage problem again takes the same symmetric form. In that environment, and under the optimal refinement characterized in the appendix, the unrestricted Bellman problem collapses to a one-dimensional choice over the probability that the next question preserves the path. This is the restricted sequence problem analyzed next.

6.3.1 Sequential Geometry: A Restricted Sequence Problem

We now represent this environment by a reduced problem in which the agent chooses directly the probability that the next question preserves the diagnostic path. We again introduce a restricted optimization problem that will serve as a tractable representation of the sequential-discipline environment.

As in the identification geometry, this problem is not obtained by rewriting the Bellman equation. Instead, it is a simplified formulation in which the agent chooses directly the probability that the next question preserves the effective path at each stage of the diagnostic process.

There is, however, an additional subtlety in the sequential environment. In the dynamic inquiry model the agent does not directly observe whether the effective path has already been broken. The agent continues to ask questions and receive answers, and the failure of the correct diagnostic trajectory is not directly identifiable from the signals themselves.

In the restricted sequence problem, by contrast, the process stops immediately once the path is broken. That is, if the agent chooses an incorrect question at any stage, the diagnostic process terminates and no further value can be obtained. At first sight this appears to impose a stronger

structure than the dynamic model.

The surprising result established below is that this restricted problem nonetheless generates the same value as the full dynamic inquiry problem. In equilibrium the optimal information policy behaves *as if* the agent immediately recognized that the path had been broken, because once success becomes impossible the marginal value of further diagnostic attention collapses to zero.

Consider a diagnostic process with $T+1$ questioning opportunities. At stage s the agent chooses a conditional success probability

$$\hat{p}_s \in \left[\frac{1}{n_s}, 1 \right],$$

where n_s denotes the number of feasible questions at that stage. The lower bound corresponds to uninformed choice among the n_s candidate questions, while the upper bound corresponds to perfectly identifying the correct next question on the effective path.

The sequence

$$(\hat{p}_T, \hat{p}_{T-1}, \dots, \hat{p}_0)$$

therefore summarizes a complete candidate strategy for the sequential diagnostic process.

Success probability. In the sequential geometry the prize is obtained only if the correct question is asked at every stage of the path. Hence the probability of ultimately obtaining the prize is

$$\Pr(\text{prize}) = \prod_{s=0}^T \hat{p}_s.$$

Define

$$\alpha_s = \prod_{r=s+1}^T \hat{p}_r,$$

the probability that stage s is reached, that is, the probability that all later questions on the path have been answered correctly.

Information costs. If conditional success probability $\hat{p}_s$ is implemented at a stage with n_s feasible candidate questions, the associated Shannon information cost is

$$C_s(\hat{p}_s) = \kappa \left[\hat{p}_s \log(n_s \hat{p}_s) + (1 - \hat{p}_s) \log\left(\frac{n_s(1 - \hat{p}_s)}{n_s - 1}\right) \right].$$

Because stage s is reached only with probability α_s , the expected information cost incurred at that stage is

$$\alpha_s C_s(\hat{p}_s).$$

The expected total information cost of the sequence is therefore

$$\sum_{s=0}^T \alpha_s C_s(\hat{p}_s).$$

The restricted optimization problem. The agent therefore chooses the sequence of conditional success probabilities to solve

$$\max_{(\hat{p}_T, \dots, \hat{p}_0)} \left\{ \prod_{s=0}^T \hat{p}_s - \sum_{s=0}^T \alpha_s C_s(\hat{p}_s) \right\}.$$

This optimization problem determines the optimal organization of inquiry when the agent directly controls the probability that each stage preserves the effective diagnostic path. In the next subsection we solve this problem explicitly.

6.3.2 Sequential Geometry: Closed-Form Solution

We now solve the restricted sequence problem introduced above.

As in the identification geometry, the agent chooses the entire sequence of conditional success probabilities

$$(\hat{p}_T, \hat{p}_{T-1}, \dots, \hat{p}_0)$$

to maximize the objective defined in the previous subsection. This is a finite-dimensional optimization problem. The choice variables are the conditional probabilities of preserving the effective path at each stage, and the objective combines the probability of remaining on the path all the way to the prize with the expected Shannon information costs incurred along the way.

To compute the solution it is convenient to exploit the dependent structure of the objective function. In the sequential geometry, the contribution of stage s matters only if all later stages have remained on the effective path. This allows the problem to be decomposed stage by stage and solved by backward induction.

As before, this recursion is a computational device for solving the restricted sequence optimization problem in closed form. The dynamic appearance arises only from the multiplicative structure of path survival in the objective.

Let $\hat{U}_s$ denote the value of the restricted sequence problem when s questioning opportunities remain. In the sequential geometry, if the effective path is preserved through all remaining stages, the prize is obtained. Thus the terminal condition is

$$\hat{U}_{-1} = 1.$$

At a stage with s periods remaining, the agent chooses the conditional success probability $p \in [1/n_s, 1]$. If the correct next question is asked, the effective path is preserved and the process continues to the next stage with value $\hat{U}_{s-1}$. If the question is incorrect, the path is broken and

continuation value falls to zero.

The stage- s problem is therefore

$$\hat{U}_s = \max_{p \in [1/n_s, 1]} \left\{ p \hat{U}_{s-1} - \kappa \left[p \log(n_s p) + (1-p) \log\left(\frac{n_s(1-p)}{n_s-1}\right) \right] \right\}.$$

The objective is strictly concave in p , so the optimum is unique. Differentiating with respect to p yields the first-order condition

$$\hat{U}_{s-1} = \kappa \log\left(\frac{p(n_s-1)}{1-p}\right).$$

Solving for the optimal conditional success probability gives

$$\hat{p}_s = \frac{1}{1 + (n_s-1)e^{-\hat{U}_{s-1}/\kappa}} = \frac{e^{\hat{U}_{s-1}/\kappa}}{e^{\hat{U}_{s-1}/\kappa} + n_s - 1}.$$

Substituting $\hat{p}_s$ back into the objective yields the recursion

$$\hat{U}_s = \kappa \log\left(\frac{e^{\hat{U}_{s-1}/\kappa} + n_s - 1}{n_s}\right).$$

This recursion can be solved explicitly. Define

$$U_s = e^{\hat{U}_s/\kappa}.$$

Then

$$U_s = \frac{U_{s-1} + n_s - 1}{n_s}, \quad U_{-1} = e^{1/\kappa}.$$

Writing

$$N_s = \prod_{r=0}^s n_r,$$

the recursion solves to

$$U_s = 1 + \frac{e^{1/\kappa} - 1}{N_s}.$$

Hence the value of the restricted sequence problem is

$$\hat{U}_s = \kappa \log\left(1 + \frac{e^{1/\kappa} - 1}{\prod_{r=0}^s n_r}\right).$$

In particular, with $T+1$ questioning opportunities, the value of the restricted sequence problem is

$$\widehat{U}_T = \kappa \log \left(1 + \frac{e^{1/\kappa} - 1}{\prod_{r=0}^T n_r} \right).$$

Thus the restricted sequence problem admits a full closed-form solution. This explicit expression will serve as the candidate value function in the verification theorem below.

6.3.3 Verification theorem

The restricted sequence problem solved above yields the explicit value function

$$\widehat{U}_T = \kappa \log \left(1 + \frac{e^{1/\kappa} - 1}{\prod_{r=0}^T n_r} \right).$$

We now verify that the candidate value function generated by the restricted sequence problem satisfies the Bellman equation of the full dynamic inquiry problem in the canonical symmetric environment.

Theorem 3 (Verification for the sequential geometry). *Consider the canonical symmetric sequential geometry in which, at each stage s , the induced diagnostic belief over the feasible questions $\mathcal{Q}_s$ is uniform and, conditional on success, the continuation problem at stage $s-1$ again takes the same symmetric form.*

Let $\widehat{U}_s$ denote the candidate value function generated by the restricted sequence problem. Then, for every stage s , the function $\widehat{U}_s$ satisfies the unrestricted Bellman recursion for the sequential geometry:

$$V_s^S(\bar{v}_s^S, \mathcal{Q}_s) = \widehat{U}_s.$$

Moreover, at each stage the supremum in the unrestricted Bellman problem is attained by the symmetric one-spike posterior family that implements the restricted success probability choice.

Hence the restricted sequence problem delivers the value of the full dynamic inquiry problem in the canonical symmetric sequential environment.

6.4 A Common Value Representation

The explicit solutions reveal that, despite their different structures, both geometries admit a common value representation. The two canonical geometries differ sharply in their diagnostic structure, but the explicit solutions derived above reveal a common underlying form. In both cases, the value of the restricted sequence problem can be written as a transformation of the probability of success under uninformed search.

Let π_T^0 denote the probability of ultimate success when, at each stage, the agent chooses uniformly among the feasible candidate questions. Then the value generated by optimal inquiry in both geometries is

$$\hat{U}_T = \kappa \log \left(1 + (e^{1/\kappa} - 1) \pi_T^0 \right).$$

The two geometries differ only in the expression for π_T^0 .

In the sequential-discipline geometry, success requires remaining on the effective path at every stage. Under uninformed search, the probability of choosing the correct question at stage r is $1/n_r$, so the baseline success probability is

$$\pi_T^0 = \frac{1}{\prod_{r=0}^T n_r}.$$

In the identification geometry, success requires locating the decisive question before the questioning opportunities are exhausted. Under uninformed search, the probability of failing at stage r is $(n_r - 1)/n_r$, so the baseline success probability is

$$\pi_T^0 = 1 - \prod_{r=0}^T \frac{n_r - 1}{n_r}.$$

Substituting these expressions into the common formula reproduces the closed-form solutions derived above. Thus the geometry of diagnosis enters only through the baseline probability of success under uninformed inquiry, while the Shannon parameter κ governs how strongly optimal attention transforms that baseline into value.

This representation separates two distinct ingredients of inquiry performance. The first is the *diagnostic geometry*, which determines how difficult the environment is absent directed inquiry. The second is the *cost of attention*, summarized by κ , which determines how effectively the agent can improve upon that baseline.

Relation to the binary Shannon rational-inattention problem. The representation

$$\hat{U}_T = \kappa \log \left(1 + (e^{1/\kappa} - 1) \pi_T^0 \right)$$

is closely related to the value formula that arises in the binary Shannon rational-inattention problem studied by Matejka and McKay (2015). Consider a binary state

$$s \in \{0, 1\}, \quad \Pr(s = 1) = \pi,$$

with payoffs $u(1) = 1$ and $u(0) = 0$. Suppose the decision maker chooses an information structure at Shannon cost $\kappa I(s; X)$. When the optimal solution is interior, the Matějka–McKay first-order conditions imply a logit state-dependent stochastic choice rule. This restriction matters: in the standard binary rational-inattention problem with asymmetric priors, the optimal policy may lie at a corner with zero attention, in which case the Matějka–McKay interior characterization does not apply and the value formula above does not obtain. Appendix E gives a simple posterior calculation for the symmetric binary Shannon problem that makes the appearance of the log–sum–exp operator

especially transparent.

Substituting this rule into the value function yields the well-known log–sum–exp representation

$$V(\pi) = \kappa \log \sum_{s \in \{0,1\}} \pi_s e^{u(s)/\kappa} = \kappa \log \left((1 - \pi) + \pi e^{1/\kappa} \right) = \kappa \log \left(1 + (e^{1/\kappa} - 1)\pi \right).$$

Evaluating this expression at $\pi = \pi_T^0$ reproduces the value formula derived above.

The similarity between the two expressions is mathematically exact, but the underlying economic problems are different. In the generic binary rational-inattention choice problem with asymmetric priors, the optimal policy may place zero attention on the state, so that the interior Matějka–McKay characterization does not apply. In the present inquiry problem, by contrast, the environment reduces to a binary margin of success—whether inquiry succeeds in locating the decisive question—and the Shannon information cost operates on the refinement of diagnostic beliefs that guide question choice.

The resulting value representation therefore uses the same Shannon log–sum–exp operator that characterizes the interior Matějka–McKay solution, even though the operator arises here from the dynamic organization of inquiry rather than from a static attention-allocation problem. The near coincidence of the formulas is therefore informative: it reflects the fact that organizing inquiry under diagnostic uncertainty transforms the probability of success through the same exponential aggregation that appears in Shannon rational-inattention models.

The common value representation also suggests a broader structural interpretation. In the environments studied here, inquiry affects payoffs only through a single terminal decision: the value of search is entirely determined by whether the decisive question is identified before the questioning opportunities are exhausted. When attention costs take the Shannon form, the resulting value expression reduces to the familiar log–sum–exp operator applied to the distribution of terminal outcomes under uninformed inquiry. This observation raises the possibility that similar representations may arise more generally in dynamic search environments where all economic consequences are realized at a single terminal decision. The present analysis does not establish such a result formally, but the coincidence suggests that the Shannon transform may provide a useful organizing device for a broader class of inquiry problems.

The Shannon transform. The common representation suggests a convenient way to summarize the value of optimal inquiry. Define the mapping

$$\mathcal{S}_\kappa(x) = \kappa \log \left(1 + (e^{1/\kappa} - 1)x \right), \quad x \in [0, 1].$$

We refer to $\mathcal{S}_\kappa(\cdot)$ as the *Shannon transform*. In both canonical diagnostic geometries the value of optimal inquiry is obtained by applying this transform to the baseline probability of success under uninformed search:

$$\widehat{U}_T = \mathcal{S}_\kappa(\pi_T^0).$$

...

The Shannon Transform of Optimal Inquiry

Let π_T^0 denote the probability of success under uninformed inquiry. Under Shannon information costs with parameter κ , optimal inquiry value satisfies

$$\hat{U}_T = \mathcal{S}_\kappa(\pi_T^0), \quad \mathcal{S}_\kappa(x) = \kappa \log\left(1 + (e^{1/\kappa} - 1)x\right).$$

Equivalently,

$$e^{\hat{U}_T/\kappa} = 1 + (e^{1/\kappa} - 1)\pi_T^0.$$

Thus optimal attention closes the gap between uninformed and fully informed value linearly on the exponential scale.

Figure 3 illustrates the Shannon transform for different information-cost parameters.

Reading the figure. Figure 3 illustrates the common value representation. The horizontal axis shows the baseline probability of success under uninformed inquiry $x = \pi_T^0$, while the vertical axis shows the value generated by optimal attention,

$$V(x) = \kappa \log\left(1 + (e^{1/\kappa} - 1)x\right).$$

The dashed 45-degree line corresponds to uninformed search, while the curves above it show how optimal attention improves upon that baseline. Smaller values of κ correspond to cheaper information acquisition and therefore to more effective inquiry

6.5 Implications and Possible Extensions

The common value representation highlights the economic role of the information-cost parameter κ . This parameter governs the marginal cost of concentrating probability on the most promising diagnostic hypothesis. Smaller values of κ correspond to cheaper information acquisition and therefore to more precise organization of inquiry.

The optimal success probabilities derived above make this role transparent. In both diagnostic geometries the conditional probability that the next question succeeds takes the form

$$p^* = \frac{1}{1 + (n - 1)e^{-\Delta/\kappa}},$$

where Δ denotes the value difference between success and failure at that stage. When κ is small, attention is cheap and the agent concentrates probability strongly on the most promising question. When κ is large, attention is costly and the agent's choice approaches uninformed selection among the candidate questions.




Figure 3: The Shannon transform $\mathcal{S}_\kappa(x)$ maps baseline success probability into the value generated by optimal inquiry. The value function

$$V(x) = \kappa \log\left(1 + (e^{1/\kappa} - 1)x\right)$$

maps the probability of success under uninformed inquiry $x = \pi_T^0$ into the value generated by optimal attention. Smaller values of κ correspond to cheaper information acquisition and therefore produce steeper improvements over the uninformed benchmark.

The representation also suggests that the structure uncovered here may extend beyond the two canonical geometries studied in this section. In both cases the value function takes the form

$$\widehat{U}_T = \kappa \log\left(1 + (e^{1/\kappa} - 1)\pi_T^0\right),$$

where π_T^0 denotes the probability of success under uninformed inquiry. This raises the possibility that similar representations may arise in a broader class of diagnostic environments, particularly in two-state exchangeable structures in which success probabilities factor multiplicatively across stages.

Finally, although the model has been presented in terms of the organization of questions, the same logic admits an equivalent interpretation in terms of the organization of listening. In that interpretation, the agent receives a stream of answers or signals and allocates attention to interpreting them correctly. The Shannon information cost then captures the difficulty of concentrating attention on the most relevant interpretation of the incoming signal.

Seen from this perspective, the model describes a more general problem of organizing inquiry and interpretation when information processing is costly. The next section turns to settings in which this process becomes observable, allowing the structure of inquiry to be measured directly.

7 Measuring Diagnostic Planning

7.1 The Historical Invisibility of Planning

Planning has always been central to economic decision-making. Workers anticipate changes in their industries, firms search for profitable strategies, and researchers attempt to identify decisive experiments. In each of these environments, outcomes depend heavily on how the decision-maker structures the inquiry that precedes action.

Yet historically the process of planning has been remarkably difficult to observe. Economists typically see the outcomes of decisions rather than the reasoning that produced them. We observe which career a worker ultimately chooses, which strategy a firm adopts, or which hypothesis a researcher tests, but the sequence of questions and answers through which those decisions were prepared is rarely visible.

This invisibility had important consequences for economic analysis. When a critical component of behavior cannot be observed directly, it becomes difficult to study systematically. As a result, economists have tended to focus on aspects of decision-making that are easier to measure—education, experience, task performance, and observable outcomes—while planning itself remained conceptually central but empirically inaccessible.

7.2 Earlier Attempts to Observe Reasoning

Researchers in psychology and decision science recognized this limitation long ago and attempted to study reasoning processes directly. The best-known approach was protocol analysis, developed most prominently by Ericsson and Simon (1980, 1993), in which participants were asked to “think aloud” while solving problems so that researchers could record the sequence of thoughts occurring during the decision process.

Protocol analysis produced valuable insights. It revealed how individuals search for solutions, interpret evidence, and revise their beliefs, and it helped establish that reasoning often proceeds through sequences of hypotheses, tests, and revisions rather than through simple optimization.

However, protocol analysis faced important limitations. Participants knew they were being observed and were asked to verbalize their reasoning explicitly. The tasks were simplified versions of real-world problems, and the stakes were typically low. As a result, the reasoning processes observed in these experiments were informative but necessarily artificial.

More fundamentally, protocol analysis could not disentangle differences in priors from differences in reasoning. If two individuals asked different questions, it was unclear whether this reflected different beliefs about the state of the world or different organization of inquiry. Planning therefore remained difficult to measure in environments where decisions actually mattered.

7.3 Observed Inquiry Records

While protocol analysis attempted to observe reasoning in laboratory settings, recent technological developments make it possible to observe inquiry directly in actual decision environments. Interactive artificial intelligence systems change this measurement constraint in a specific way: inquiry becomes explicit, sequential, and preserved.

When individuals interact with AI systems while exploring decisions, the sequence of questions and answers that precedes a decision is recorded. Each prompt generates a response, and each response may lead to new questions, revisions, or shifts in direction. These records capture the structure of inquiry in a way that was previously impossible.

Several features make these data particularly valuable. They arise in actual decision contexts rather than artificial laboratory tasks, they are inherently sequential—each question appearing in the context created by previous answers—and they are incentivized, since users ask questions in pursuit of real objectives such as employment, strategy, or clarification of uncertainty. Taken together, these characteristics create an unprecedented opportunity to observe planning as it actually unfolds.

From model objects to observable data. The model developed above describes planning as a sequential process of inquiry. At each stage the decision-maker selects a question q_t , receives an answer a_t , and updates beliefs before selecting the next question. Formally, inquiry is a stochastic process over sequences

$$(q_1, a_1), (q_2, a_2), \dots, (q_T, a_T).$$

Historically these sequences were rarely observable. Decisions were recorded, but the inquiries that preceded them were not. Interactive AI systems change this constraint by preserving the sequence of prompts and responses through which decisions are explored. We refer to these sequences as *inquiry paths*.

Inquiry Path

An inquiry path is the sequential record of inquiry produced during decision exploration:

$$\mathcal{P} = \{(q_1, a_1), (q_2, a_2), \dots, (q_T, a_T)\}.$$

Each element records a question posed by the agent and the answer generated by the environment. Inquiry paths preserve the observable structure of inquiry: the ordering of questions, the branching of investigation, and the point at which inquiry converges or stops.

Inquiry paths reveal ordering, branching, and the points at which investigation stabilizes or terminates. What they do not directly reveal are the internal beliefs or refinement processes that generated them. Interpreting these records therefore requires a model that links observable question sequences to latent diagnostic beliefs and refinement decisions. Without such a model, they are

descriptive; with such a model, they become observable realizations of structured inquiry.i

7.4 From Records to Measurement

Observed inquiry records capture behavior, but they do not by themselves measure planning skill. Their interpretation requires a theory of how inquiry should be organized.

The framework developed in this paper provides such a structure. In the model, planning corresponds to the sequencing of questions when the diagnostic value of those questions is uncertain. The agent chooses questions according to a stochastic rule that depends on beliefs about diagnostic structure, and answers update those beliefs through Bayesian updating. An inquiry path

$$\mathcal{P} = \{(q_1, a_1), \dots, (q_T, a_T)\}$$

can therefore be interpreted as a realized path of the dynamic inquiry process.

Natural data provide rich observational variation but embed uncontrolled heterogeneity in diagnostic structure and priors. The model therefore suggests a complementary experimental approach.

7.5 Observed Data and Structured Experiments

The framework developed in this paper provides a structure for interpreting observed inquiry records. In the model, planning corresponds to the sequencing of questions when the diagnostic value of those questions is uncertain.

In structured environments the diagnostic regime D can be drawn from a known distribution and the experiment menu $\{\Pi_q^D\}$ explicitly defined. Participants ask questions from a finite set $\mathcal{Q}$ and receive answers generated by the underlying diagnostic structure.

In these environments the sequence of questions chosen by the participant directly determines expected success. Failure to identify decisive distinctions lowers expected payoff, while excessive exploration of irrelevant branches consumes the finite inquiry budget. Refinement intensity becomes a structural parameter that can be estimated from observed inquiry paths.

7.6 Structured Life-Event Tasks

Laboratory precision alone does not capture the stakes of real planning environments.

Many consequential life events already possess the structure required by the model.

Consider the construction of a CV for a specific job. The participant begins with a large informational base: prior roles, projects, credentials, and experience. A job description implicitly defines the diagnostic regime. Some distinctions are decisive while others are irrelevant.

Interacting with an AI system, the participant must decide which distinctions to foreground, which uncertainties to clarify, and when further inquiry is unnecessary.

Two participants facing the same task may organize inquiry very differently. Some quickly identify decisive distinctions. Others explore descriptive features of the environment without isolating the structural constraint that determines success.

These differences correspond precisely to the planning capabilities described in the model.

7.7 Co-Evolution of Model and Measurement

Historically the absence of durable inquiry records constrained the study of planning. Planning before commitment was conceptually central but empirically inaccessible.

Interactive AI systems remove that constraint. Inquiry becomes explicit, sequential, and preserved.

Measurement and modeling therefore co-evolve. The formal architecture developed in this paper provides the structure through which preserved inquiry paths can be interpreted. Conversely, the existence of such records makes it possible to measure planning behavior directly. Planning under diagnostic uncertainty thus becomes a measurable object of economic analysis.

The emergence of durable inquiry records is closely tied to the spread of interactive artificial intelligence systems. The next section examines how these systems reshape the economic environment in which planning occurs.

8 Planning in the Age of Artificial Intelligence

Up to this point the analysis has been deliberately technology-neutral. The model describes how agents organize inquiry when the diagnostic value of questions is uncertain. Artificial intelligence alters the economic environment precisely along this margin.

Modern AI systems can generate candidate answers to a wide range of questions almost instantly. Information that once required substantial search or consultation can now be produced within seconds. As a result, the central constraint in decision-making shifts from acquiring answers to organizing inquiry. The bottleneck moves from information acquisition to the sequencing of questions.

Artificial intelligence therefore places the organization of inquiry at the center of economic decision-making.

Planning Complementarity Conjecture.

Artificial intelligence dramatically reduces the cost of generating answers to questions. As answer generation becomes inexpensive, the binding constraint in decision-making shifts from information acquisition to the organization of inquiry.

Planning Complementarity Conjecture: reductions in the cost of producing answers increase the economic return to planning skill, defined as the ability to generate decisive questions and organize sequences of inquiry that isolate the relevant distinctions in a decision problem.

Under this conjecture, advances in artificial intelligence do not eliminate the role of human judgment. Instead they raise the payoff to disciplined inquiry by making differences in how individuals organize questions more consequential for outcomes.

8.1 Revisiting the Inquiry Problem

Consider again a worker exploring whether a transition into data science is feasible. Using an artificial intelligence system, the worker can generate answers to many descriptive questions almost instantly:

- What jobs exist in data science?
- What salaries do these roles offer?
- Which cities have the largest job markets?
- What programming languages are commonly required?

Each answer provides useful descriptive information, yet the investigation may expand indefinitely without resolving the central decision. A different sequence of questions produces a different structure of inquiry:

- What determines whether someone can enter the field?
- Is a specific credential required?
- Given my background, how difficult would it be to obtain it?
- What is the fastest credible way to do so?

In this sequence the investigation converges rapidly on the decisive constraint. In the language of the model, the two sequences correspond to different diagnostic strategies: the first explores descriptive dimensions of the environment, while the second isolates the decisive diagnostic distinction. Artificial intelligence does not remove the need for planning; it makes the consequences of how inquiry is organized more visible.

8.2 Planning Skill

When answers become abundant, differences in planning skill become more visible. In the model, planning skill manifests through several capabilities in organizing inquiry.

1. **Question generation.** Individuals differ in the questions they are able to formulate. Some planners quickly generate questions that expose decisive distinctions, while others explore descriptive features of the environment without identifying the structural pivot of the problem.
2. **Answer interpretation.** Answers must be interpreted in ways that redirect investigation. Extracting the relevant implication from a response often requires judgment about how the information reshapes the decision problem.
3. **Anticipation of inquiry paths.** Skilled planners anticipate which future questions are likely to be informative and direct inquiry toward those distinctions.

Artificial intelligence does not eliminate these differences. It exposes and amplifies them.

8.3 Complementarity and Substitution

The interaction between artificial intelligence and planning skill depends fundamentally on the diagnostic structure of the environment. In relatively simple environments, a small number of well-chosen questions can resolve most of the relevant uncertainty. When answers to those questions can be generated quickly and reliably, artificial intelligence has the potential to substitute for routine forms of reasoning. The bottleneck is limited, and reducing the cost of answers directly improves performance.

The situation changes in diagnostically complex environments. In such settings, questions do not operate independently: early answers reshape the space of possibilities and alter which subsequent questions are informative. Inquiry becomes path-dependent, and the value of any given question depends on the sequence in which it is asked. Under these conditions, the ability to identify decisive distinctions and to organize sequences of questions becomes central to performance.

Artificial intelligence therefore changes not only the level of productivity, but the relative importance of different cognitive skills. When answers are abundant, differences in how individuals structure inquiry become more consequential. Individuals with similar access to technology may nevertheless achieve very different outcomes, depending on their planning discipline and their ability to allocate attention to the most informative lines of inquiry.

The framework developed in this paper provides a way to analyze this interaction formally. Differences in planning discipline correspond to differences in the question sets agents consider, the beliefs they hold about the diagnostic value of those questions, and the cognitive resources they allocate to evaluating and selecting among them. Artificial intelligence amplifies the returns to these choices by reducing the cost of obtaining answers once a question has been posed.

8.4 Relation to Task-Based Frameworks

The dominant empirical framework for analyzing technological change in labor markets focuses on tasks. In this approach, production is decomposed into activities, and technological progress is studied in terms of how it substitutes for or complements the execution of those activities (Autor et al., 2003; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2020). This perspective has been highly successful in organizing evidence on automation and labor demand.

The framework developed here highlights a different margin of adjustment. Rather than focusing on the execution of tasks, it emphasizes how agents organize inquiry before committing to action. From this perspective, the central object is not the task itself, but the process through which the agent determines what to do.

Artificial intelligence systems make this distinction particularly salient. These systems do not simply automate tasks; they generate answers to questions. As the cost of obtaining answers falls, the productivity of decision-making increasingly depends on how effectively agents structure their inquiry. The key constraint shifts from execution to diagnosis.

Planning under diagnostic uncertainty therefore operates at a different level of analysis than the task-based framework. The two perspectives are complementary. Task-based models describe

how activities are performed, while the framework developed here describes how those activities are selected and organized in the first place.

8.5 Measurement and Development

Historically, the organization of inquiry has been extremely difficult to observe. Economists typically see the outcomes of decisions rather than the process through which those decisions were prepared. As a result, planning has remained conceptually central but empirically elusive.

Artificial intelligence systems fundamentally change this constraint. Interaction logs preserve the sequence of questions through which decisions are explored, providing a direct record of how individuals structure inquiry under uncertainty. What was previously hidden becomes observable: the questions that are asked, the order in which they are posed, and the responses that guide subsequent investigation.

When the diagnostic structure of an environment can be specified in advance, these inquiry sequences can be evaluated objectively. Questions differ in their informational value, and some identify decisive distinctions while others pursue irrelevant lines of inquiry. Participants who focus on high-value distinctions increase the probability of reaching the correct decision, while those who do not reduce it. Performance can therefore be assessed not only from final outcomes, but from the structure of the inquiry that produced them.

The same environments that permit objective assessment also create opportunities for deliberate practice. Exercises built around hidden constraints, source reliability, or convergence decisions allow individuals to develop the ability to identify decisive distinctions and to organize inquiry efficiently. Because artificial intelligence systems both generate answers and record interaction histories, they provide a natural platform for constructing and scaling such environments.

Planning capital therefore becomes not only measurable, but teachable. The organization of inquiry can be observed, evaluated, and systematically improved in ways that were previously infeasible.

From observable inquiry to economic analysis When inquiry leaves a durable record, the organization of investigation becomes observable rather than inferred. Differences in how agents generate questions, interpret answers, and sequence inquiry can therefore be studied directly, rather than reconstructed indirectly from outcomes.

The framework developed in this paper provides a structure for interpreting these observations. It identifies the economic object underlying inquiry traces as planning under diagnostic uncertainty, and provides a language for linking observed question sequences to underlying beliefs, constraints, and strategies.

These developments suggest a shift in what can be measured in economic analysis. The organization of inquiry need no longer be treated as an unobserved residual. Instead, it can be incorporated directly as a measurable component of behavior, opening the door to new forms of theoretical modeling, empirical analysis, and policy design.

9 Intellectual Foundations and Related Traditions

The framework developed in this paper extends the Cognitive Economics research program, which studies how cognitive processes shape economic behavior and therefore require explicit modeling and measurement. Within this perspective, attention allocation, belief formation, and the generation of mistakes are treated as economically consequential objects rather than residuals. A central methodological implication of this approach is data engineering: the design of environments and data structures that render cognitive processes observable in economically meaningful settings (Caplin, 2025b, 2023, 2025a).

From this perspective, the absence of a formal model of inquiry is not incidental. In standard economic data, we observe outcomes—choices, earnings, employment transitions—but we do not observe the process through which decision-makers generate the information that leads to those outcomes. The sequence of questions that are asked, the sources that are consulted, and the way attention is allocated across them are typically unobserved. As a result, planning remains implicit: models assume that relevant information is available, but do not describe how it is produced. The present paper contributes to this agenda by introducing a formal model in which sequential question choice under diagnostic uncertainty is itself a constrained economic decision, and by providing a structure that links this object to observable inquiry behavior.

9.1 Inquiry and the Selection of Relevant Distinctions

In many decision problems, the central difficulty lies not only in how agents respond to information, but in determining which information is relevant in the first place. Decision-makers must decide which distinctions to investigate, which sources to consult, and how to allocate attention across them. These choices shape what is learned before any final action is taken.

This aspect of decision-making is often implicit in economic models. Agents are assumed to face a set of alternatives and an informational environment that is already structured in a way that makes evaluation possible. The problem is then to choose optimally given that structure. In practice, however, the structure of the problem is not fully determined in advance. It is shaped by the process through which information is generated and interpreted.

Herbert Simon was distinctive in emphasizing this layer. In his work on problem solving and bounded rationality (Simon, 1955, 1996; Newell and Simon, 1972), decision-making is not only a matter of choosing among alternatives, but of identifying which distinctions matter and how complex environments are decomposed into manageable components. This perspective highlights that the effectiveness of decision-making depends on how attention is directed and how inquiry is organized.

At the same time, this literature remains largely descriptive. It identifies the importance of structuring problems and directing inquiry, but does not provide a formal account of how agents determine which distinctions are informative or how they allocate resources to learning about them under constraints.

The present paper takes a step in this direction. It treats the selection of informative distinctions as an economic decision problem in its own right. Agents face uncertainty not only about payoff-relevant states, but about how inquiry reveals those states, and must allocate attention to refining beliefs about which questions are informative. This provides a formal framework in which the organization of inquiry can be analyzed and, in principle, measured.

9.2 Information acquisition under payoff-state uncertainty

Economic theory provides several rigorous frameworks for analyzing learning before commitment. Search models study how agents sample from known distributions and decide when to stop searching (Mortensen and Pissarides, 1994; Rogerson et al., 2005). Bandit models analyze exploration across a set of known experiments, balancing the value of information against immediate rewards (Gittins, 1979; Berry and Fristedt, 1985). Rational inattention models characterize how agents allocate limited information-processing capacity across signals about payoff-relevant states (Sims, 2003; Matějka and McKay, 2015; Caplin and Dean, 2015; Caplin et al., 2022; Woodford, 2014).

These frameworks differ in their primitives and in the margins along which optimization occurs, but they share a common structural feature: uncertainty is located at the level of payoff-relevant states. The informational environment—how signals are generated from those states—is specified *ex ante*. Agents may choose how intensively to search, which experiment to sample, or how much attention to allocate, but the mapping from states to signals, or from actions to observations, is taken as known. Learning therefore concerns the payoff-relevant state, not the structure through which information about that state is produced.

This feature has been central to the tractability and success of these models, but it also limits their scope. When the informational structure is given, the problem of determining which inquiries are informative does not arise. The present paper shifts the location of uncertainty from the payoff-relevant state to the diagnostic structure itself. Agents are uncertain not only about θ , but about how inquiry reveals θ , and may allocate costly attention to refining beliefs about which questions are informative. The resulting problem is therefore not simply one of learning about the state, but of learning how information about the state is generated. This shift introduces a higher-order learning problem over the structure of information itself.

9.3 Planning heterogeneity, decision skill, and measurement

A large empirical literature documents heterogeneity in cognitive and non-cognitive skills, including decision-making ability, foresight, and planning behavior. Early work emphasized the role of planning in financial behavior, showing that individuals differ systematically in their propensity to plan and that these differences are associated with economically meaningful outcomes (Ameriks et al., 2003). More recent work has expanded the scope of measurement to include broader decision-making skills and their relationship to income and life outcomes (Caplin et al., 2023), as well as the role of cognitive and non-cognitive abilities in shaping labor market performance (Heckman et al., 2006; Deming, 2022).

At the same time, research in psychology has studied planning ability directly, often using structured tasks that require sequential reasoning and foresight (Ericsson and Simon, 1980). These studies demonstrate substantial variation in the ability to organize multi-step reasoning and adapt strategies over time. Related work emphasizes the forward-looking structure of planning, including the formation of intentions under anticipated contingencies (Ajzen, 1991) and the use of implementation intentions to link specific situations to actions (Gollwitzer, 1999). Together, these approaches highlight that planning involves not only reasoning about outcomes, but structuring responses to future states of the world. However, such measures are typically derived from laboratory environments or stylized tasks, and are only indirectly connected to economically consequential decisions. This perspective aligns with the view taken here, in which planning concerns how agents organize responses to uncertain and evolving informational environments.

A common feature of both literatures is that planning is either measured in reduced form or inferred from outcomes. The underlying process by which agents structure inquiry is rarely observed directly or modeled explicitly. The framework developed in this paper provides a bridge between these approaches by introducing a formal model of sequential question choice under diagnostic uncertainty and by linking it to observable inquiry traces. This makes it possible to study planning as a structured economic object rather than as a residual category or latent trait.

9.4 Artificial intelligence and the observability of inquiry

Recent advances in artificial intelligence have fundamentally altered the environment in which inquiry takes place. Large language models and related systems can generate candidate answers to a wide range of questions at very low cost (OpenAI, 2023). Empirical evidence suggests that such systems can substantially increase productivity in a variety of tasks (Noy and Zhang, 2023; Brynjolfsson et al., 2025). Much of the existing economic analysis has focused on how these technologies affect the execution of tasks and the allocation of labor across activities.

The framework developed here highlights a different margin. As the cost of generating answers declines, the constraint on decision-making shifts toward the organization of inquiry. When answers are abundant, the key question becomes which questions to ask and in what order. Differences in the ability to structure inquiry may therefore play an increasingly important role in shaping economic outcomes, even when access to technology is similar.

A second implication concerns measurement. Interactions with AI systems naturally generate sequential records of inquiry, including the questions asked, the answers received, and the points at which inquiry stops. These traces provide a new form of data on the structure of reasoning in economically relevant settings. The model developed in this paper interprets such traces as realizations of a dynamic inquiry process under diagnostic uncertainty, thereby linking advances in AI to the empirical study of planning behavior.

Taken together, these developments underscore the importance of modeling inquiry itself as an economic object. The present framework identifies planning under diagnostic uncertainty as a distinct margin of behavior that becomes both more consequential and more observable in environ-

ments where answers are inexpensive but attention and structure remain scarce.

10 Conclusion

Economic decisions are rarely made in a single step. They are preceded by inquiry. Workers exploring career transitions, firms evaluating strategies, and researchers designing empirical tests all engage in sequential investigation before committing to action. In each case, the quality of the final decision depends not only on the information available but on how the inquiry preceding action is organized.

Historically this margin remained largely invisible. Economists observed the outcomes of decisions rather than the reasoning that produced them, and the sequence of questions through which decisions were prepared rarely left a durable trace. Planning was therefore conceptually central but empirically inaccessible.

Artificial intelligence changes this constraint. When answers become inexpensive to generate, the scarce input in decision-making shifts from acquiring information to organizing inquiry. At the same time, inquiry leaves a record: the sequence of questions and answers through which decisions are explored becomes observable.

The framework developed in this paper formalizes this margin. Planning under diagnostic uncertainty becomes a defined economic object rather than an informal intuition. Once observable, it can also be measured and developed. The organization of inquiry thus becomes a measurable and economically consequential component of decision-making in environments where answers are abundant but understanding remains scarce.

References

- Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. *Handbook of Labor Economics*, 4B:1043–1171.
- Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, 128(6):2188–2244.
- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2):179–211.
- Ameriks, J., Caplin, A., and Leahy, J. (2003). Wealth accumulation and the propensity to plan. *Quarterly Journal of Economics*, 118(3):1007–1047.
- Autor, D., Levy, F., and Murnane, R. (2003). The skill content of recent technological change. *Quarterly Journal of Economics*, 118(4):1279–1333.
- Berry, D. A. and Fristedt, B. (1985). *Bandit Problems: Sequential Allocation of Experiments*. Chapman and Hall.

- Brynjolfsson, E., Li, D., and Raymond, L. (2025). Generative ai at work. *The Quarterly Journal of Economics*, 140(2):889–942.
- Brynjolfsson, E., Li, D., and Raymond, L. R. (2023). Generative ai at work. *Quarterly Journal of Economics*. NBER Working Paper 31161; forthcoming in QJE.
- Caplin, A. (2023). *Science Of Mistakes, The: Lecture Notes On Economic Data Engineering*, volume 16. World Scientific.
- Caplin, A. (2025a). Data engineering for cognitive economics. *Journal of Economic Literature*, 63(1):164–196.
- Caplin, A. (2025b). *An introduction to cognitive economics: The science of mistakes*. Springer Nature.
- Caplin, A. and Dean, M. (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7):2183–2203.
- Caplin, A., Dean, M., and Leahy, J. (2022). Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political Economy*, 130(6):1676–1715.
- Caplin, A., Deming, D. J., Leth-Petersen, S., and Weidmann, B. (2023). Economic decision-making skill predicts income in two countries. Technical report, National Bureau of Economic Research.
- Deming, D. (2022). Four facts about human capital. *Journal of Economic Perspectives*, 36(3):75–102.
- Ericsson, K. A. and Simon, H. A. (1980). Verbal reports as data. *Psychological Review*, 87(3):215–251.
- Ericsson, K. A. and Simon, H. A. (1993). *Protocol Analysis: Verbal Reports as Data*. MIT Press.
- Gittins, J. C. (1979). *Bandit Processes and Dynamic Allocation Indices*. Journal of the Royal Statistical Society Series B (Supplement).
- Gollwitzer, P. M. (1999). Implementation intentions: Strong effects of simple plans. *American Psychologist*, 54(7):493–503.
- Heckman, J., Stixrud, J., and Urzua, S. (2006). The effects of cognitive and noncognitive abilities on labor market outcomes. *Journal of Labor Economics*, 24(3):411–482.
- Matejka, F. and McKay, A. (2015). Rational inattention to discrete choices. *American Economic Review*, 105(1):272–298.
- Matějka, F. and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–298.

- Mortensen, D. T. and Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *Review of Economic Studies*, 61(3):397–415.
- Newell, A. and Simon, H. A. (1972). *Human Problem Solving*. Prentice-Hall, Englewood Cliffs, NJ.
- Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192.
- OpenAI (2023). Gpt-4 technical report. Technical report, OpenAI.
- Rogerson, R., Shimer, R., and Wright, R. (2005). Search-theoretic models of the labor market: A survey. *Journal of Economic Literature*, 43(4):959–988.
- Savage, L. J. (1954). *The Foundations of Statistics*. John Wiley & Sons, New York.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1):99–118.
- Simon, H. A. (1996). *The Sciences of the Artificial*. MIT Press, Cambridge, MA, 3 edition.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50:665–690.
- Woodford, M. (2014). Rational inattention and state dependent stochastic choice. *Journal of Economic Theory*, 153:534–558.

A Structure of Appendix

This appendix contains two types of results. The first set of sections provides verification proofs for the main theorems. The remaining sections provide auxiliary calculations that clarify the closed-form solutions and their connection to Shannon rational-inattention models.

B Proof of Theorem 1

We prove existence and equivalence by backward induction on the horizon.

B.1 Preliminaries

Fix a period t and an admissible state $(\mu_t, \mathcal{Q}_t)$. Because Θ , $\mathcal{D}$, and $\mathcal{Q}_t$ are finite, the SDSC domain

$$\mathcal{R}(\mathcal{Q}_t) = \prod_{(\theta, D) \in \Theta \times \mathcal{D}} \Delta(\mathcal{Q}_t)$$

is a nonempty compact subset of a finite-dimensional Euclidean space.

The immediate-commitment value

$$\bar{V}(\mu_t) = \max_{y \in \mathcal{Y}} \sum_{\theta} \mu_t^{\Theta}(\theta) U(y, \theta)$$

is finite for every μ_t , since $\mathcal{Y}$ and Θ are finite and utility is bounded.

Let

$$\nu_t(D) = \sum_{\theta} \mu_t(\theta, D)$$

denote the marginal belief over diagnostic structures.

In the posterior formulation, continuation from $(\mu_t, \mathcal{Q}_t)$ is governed by refinements $\Gamma \in \mathcal{B}(\nu_t)$ together with choice rules $\alpha(\cdot \mid \gamma)$ defined on refined beliefs

$$\gamma(\theta, D) = \mu_t(\theta \mid D) \hat{\nu}(D), \quad \hat{\nu} \sim \Gamma.$$

In the SDSC formulation, continuation is governed by rules $\rho_t \in \mathcal{R}(\mathcal{Q}_t)$ together with the induced cost

$$C^S(\rho_t; \mu_t, \mathcal{Q}_t) = \inf \{ \mathcal{C}(\nu_t, \Gamma) : \rho_t \text{ is consistent with } \Gamma \}.$$

By Assumption 1, for every $(\mu_t, \mathcal{Q}_t)$ this cost is finite, the infimum is attained, and $C^S(\cdot; \mu_t, \mathcal{Q}_t)$ is continuous on $\mathcal{R}(\mathcal{Q}_t)$.

B.2 Step 1: Existence for the SDSC problem

We first prove existence of an optimal SDSC rule at every state.

For $t = -1$, there is no continuation problem and

$$V_{-1}^S(\mu, \mathcal{Q}) = \bar{V}(\mu),$$

so existence is immediate.

Now fix $t \geq 0$ and suppose inductively that $V_{t-1}^S(\mu, \mathcal{Q})$ is finite for every admissible $(\mu, \mathcal{Q})$.

Consider the SDSC continuation problem at $(\mu_t, \mathcal{Q}_t)$:

$$J_t^S(\mu_t, \mathcal{Q}_t) = \max_{\rho_t \in \mathcal{R}(\mathcal{Q}_t)} \left\{ -C^S(\rho_t; \mu_t, \mathcal{Q}_t) + \sum_{q \in \mathcal{Q}_t} \sum_{a \in A_q} \mathbb{P}_t(q, a) V_{t-1}^S(B_{(\mu_t, q, a)}, \mathcal{Q}_t \setminus \{q\}) \right\},$$

where the induced joint distribution over questions and answers is

$$\mathbb{P}_t(q, a) = \sum_{\theta, D} \mu_t(\theta, D) \rho_t(q \mid \theta, D) P(a \mid \theta, D, q).$$

Because Θ , $\mathcal{D}$, and $\mathcal{Q}_t$ are finite, the map

$$\rho_t \mapsto \mathbb{P}_t(q, a)$$

is continuous for each (q, a) , and hence the objective function is continuous in ρ_t . By Assumption 1, $C^S(\rho_t; \mu_t, \mathcal{Q}_t)$ is continuous. Therefore the entire objective is continuous on the compact set $\mathcal{R}(\mathcal{Q}_t)$.

By the Weierstrass theorem, a maximizer

$$\rho_t^* \in \mathcal{R}(\mathcal{Q}_t)$$

exists. It follows that $J_t^S(\mu_t, \mathcal{Q}_t)$ is well defined and finite, and hence so is

$$V_t^S(\mu_t, \mathcal{Q}_t) = \max \{ \bar{V}(\mu_t), J_t^S(\mu_t, \mathcal{Q}_t) \}.$$

This proves part (1) of the theorem.

B.3 Step 2: One-step equivalence of continuation values

We now show that, at any given state $(\mu_t, \mathcal{Q}_t)$, the posterior/refinement and SDSC continuation problems have the same value.

Posterior formulation. The continuation value is

$$J_t(\mu_t, \mathcal{Q}_t) = \sup_{\Gamma \in \mathcal{B}(\nu_t)} \left[-\mathcal{C}(\nu_t, \Gamma) + \mathbb{E}_{\hat{\nu} \sim \Gamma} \left(\max_{q \in \mathcal{Q}_t} \sum_{a \in A_q} P(a \mid \gamma, q) V_{t-1}(B_{(\gamma, q, a)}, \mathcal{Q}_t \setminus \{q\}) \right) \right],$$

where $\gamma(\theta, D) = \mu_t(\theta \mid D) \hat{\nu}(D)$ and

$$P(a \mid \gamma, q) = \sum_{\theta, D} \gamma(\theta, D) P(a \mid \theta, D, q).$$

SDSC formulation. The continuation value is

$$J_t^S(\mu_t, \mathcal{Q}_t) = \max_{\rho_t \in \mathcal{R}(\mathcal{Q}_t)} \left\{ -C^S(\rho_t; \mu_t, \mathcal{Q}_t) + \sum_{q \in \mathcal{Q}_t} \sum_{a \in A_q} \mathbb{P}_t(q, a) V_{t-1}^S(B_{(\mu_t, q, a)}, \mathcal{Q}_t \setminus \{q\}) \right\}.$$

B.3.1 Posterior value is no greater than SDSC value

Fix any feasible (Γ, α) in the posterior problem and let ρ_t be the induced SDSC rule.

Then (Γ, α) and ρ_t generate the same joint law over (θ, D, q, a) and hence the same distribution over next-period beliefs. Therefore

$$\begin{aligned} & \mathbb{E}_{\hat{\nu} \sim \Gamma} \left[\sum_{q \in \mathcal{Q}_t} \alpha(q \mid \gamma) \sum_{a \in A_q} P(a \mid \gamma, q) V_{t-1}^S(B_{(\gamma, q, a)}, \mathcal{Q}_t \setminus \{q\}) \right] \\ &= \sum_{q \in \mathcal{Q}_t} \sum_{a \in A_q} \mathbb{P}_t(q, a) V_{t-1}^S(B_{(\mu_t, q, a)}, \mathcal{Q}_t \setminus \{q\}). \end{aligned}$$

Since ρ_t is consistent with Γ ,

$$C^S(\rho_t; \mu_t, \mathcal{Q}_t) \leq \mathcal{C}(\nu_t, \Gamma).$$

Using the inductive hypothesis $V_{t-1}^S = V_{t-1}$, it follows that

$$J_t(\mu_t, \mathcal{Q}_t) \leq J_t^S(\mu_t, \mathcal{Q}_t).$$

B.3.2 SDSC value is no greater than posterior value

Fix any $\rho_t \in \mathcal{R}(\mathcal{Q}_t)$. By Assumption 1, there exists (Γ^*, α^*) consistent with ρ_t such that

$$C^S(\rho_t; \mu_t, \mathcal{Q}_t) = \mathcal{C}(\nu_t, \Gamma^*).$$

Again, (Γ^*, α^*) and ρ_t generate the same joint law over (θ, D, q, a) and hence

$$\sum_{q \in \mathcal{Q}_t} \sum_{a \in A_q} \mathbb{P}_t(q, a) V_{t-1}^S(B_{(\mu_t, q, a)}, \mathcal{Q}_t \setminus \{q\}) = \mathbb{E}_{\hat{\nu} \sim \Gamma^*} \left[\sum_{q \in \mathcal{Q}_t} \alpha^*(q \mid \gamma) \sum_{a \in A_q} P(a \mid \gamma, q) V_{t-1}(B_{(\gamma, q, a)}, \mathcal{Q}_t \setminus \{q\}) \right].$$

Thus ρ_t attains a value feasible in the posterior problem, implying

$$J_t^S(\mu_t, \mathcal{Q}_t) \leq J_t(\mu_t, \mathcal{Q}_t).$$

Combining both inequalities yields

$$J_t(\mu_t, \mathcal{Q}_t) = J_t^S(\mu_t, \mathcal{Q}_t).$$

B.4 Step 3: Equality of value functions

Since $J_t = J_t^S$ and both value functions are defined by the same maximum with $\bar{V}(\mu_t)$, it follows that

$$V_t(\mu_t, \mathcal{Q}_t) = V_t^S(\mu_t, \mathcal{Q}_t).$$

B.5 Conclusion

By backward induction, for every horizon t and every admissible state $(\mu_t, \mathcal{Q}_t)$, the SDSC and posterior formulations admit optimal policies and yield identical value functions. This proves Theorem 1.

C Verification for Identification and Sequential Geometries

This appendix verifies that the explicit value functions derived from the restricted sequence problems satisfy the Bellman equations for the two diagnostic geometries. In both cases the proof

proceeds by induction on the horizon and uses the one-period Shannon formula for the corresponding reduced choice problem.

C.1 Identification geometry

This appendix proves the verification result for the identification geometry in the canonical symmetric environment.

Let $\mathcal{Q}_s$ denote the feasible question set at stage s , and let $J_s := |\mathcal{Q}_s|$. In the canonical symmetric environment, the induced diagnostic belief over the decisive question is uniform:

$$\bar{\nu}_s^I(q) = \frac{1}{J_s}, \quad q \in \mathcal{Q}_s.$$

Let $V_s^I(\nu, \mathcal{Q}_s)$ denote the value of the unrestricted Bellman problem, and let $\hat{U}_s^I(J_s)$ denote the candidate value generated by the restricted sequence problem.

Theorem 2 (Verification for the identification geometry). Consider the canonical symmetric identification geometry in which, at each stage s , the induced diagnostic belief over the currently feasible questions $\mathcal{Q}_s$ is uniform.

Let $\hat{U}_T$ denote the value generated by the restricted sequence problem, and let $V_T^I(\nu_T^I, \mathcal{Q}_T)$ denote the value of the unrestricted Bellman problem for the identification geometry.

Then the restricted value coincides with the unrestricted value in the canonical symmetric environment:

$$V_T^I(\bar{\nu}_T^I, \mathcal{Q}_T) = \hat{U}_T.$$

Equivalently, the symmetric one-spike posterior family that implements the restricted sequence problem attains the supremum in the unrestricted Bellman problem.

Proof. The proof proceeds by backward induction on the number of remaining inquiry opportunities.

Step 1: Final period. When $s = 0$, the agent has one final opportunity to refine beliefs and ask a question. Since no continuation remains after failure, the unrestricted Bellman problem is

$$V_0^I(\bar{\nu}_0^I, \mathcal{Q}_0) = \sup_{\Gamma_0 \in \mathcal{B}(\bar{\nu}_0^I)} \left[-\kappa I(\Gamma_0; \bar{\nu}_0^I) + \mathbb{E}_{\hat{\nu}_0 \sim \Gamma_0} \left(\max_{q \in \mathcal{Q}_0} \hat{\nu}_0(q) \right) \right].$$

This is a one-period Shannon rational-inattention problem with J_0 ex ante symmetric actions. By the Matějka–McKay characterization, the unrestricted optimum is implemented by a symmetric one-spike posterior family and yields

$$V_0^I(\bar{\nu}_0^I, \mathcal{Q}_0) = \kappa \log \left(\frac{e^{1/\kappa} + J_0 - 1}{J_0} \right).$$

Equivalently, the optimal probability that the final question is decisive is

$$\hat{p}_0 = \frac{e^{1/\kappa}}{e^{1/\kappa} + J_0 - 1}.$$

This is exactly the terminal value generated by the restricted sequence problem, so

$$V_0^I(\bar{\nu}_0^I, \mathcal{Q}_0) = \hat{U}_0^I(J_0).$$

Step 2: Inductive hypothesis. Fix $s \geq 1$, and assume that at stage $s - 1$ the unrestricted Bellman problem coincides with the candidate restricted value:

$$V_{s-1}^I(\bar{\nu}_{s-1}^I, \mathcal{Q}_{s-1}) = \hat{U}_{s-1}^I(J_{s-1}).$$

Step 3: Candidate posterior family. Consider the symmetric one-spike posterior family indexed by a scalar success probability $p \in [1/J_s, 1]$. For each $q \in \mathcal{Q}_s$, define the posterior $\hat{\nu}^{q,p}$ by

$$\hat{\nu}^{q,p}(q) = p, \quad \hat{\nu}^{q,p}(q') = \frac{1-p}{J_s-1}, \quad q' \neq q.$$

Let these posteriors occur with equal ex ante probability $1/J_s$. This family is Bayes-plausible under the uniform prior $\bar{\nu}_s^I$.

If the agent chooses the posterior-maximizing question q and that question fails, then the renormalized posterior over the remaining questions is uniform:

$$(\hat{\nu}^{q,p})^{-q}(q') = \frac{1}{J_s-1}, \quad q' \in \mathcal{Q}_s \setminus \{q\}.$$

Hence failure returns the problem to the canonical symmetric environment one stage earlier.

Step 4: Guess and verify. Under the candidate posterior family, choosing the posterior-maximizing question yields current payoff

$$p + (1-p)\hat{U}_{s-1}^I(J_s - 1).$$

The Shannon information cost of implementing the one-spike family is

$$\kappa \left[p \log(J_s p) + (1-p) \log\left(\frac{J_s(1-p)}{J_s-1}\right) \right].$$

Therefore the value generated by the candidate family is

$$p + (1-p)\hat{U}_{s-1}^I(J_s - 1) - \kappa \left[p \log(J_s p) + (1-p) \log\left(\frac{J_s(1-p)}{J_s-1}\right) \right].$$

We now verify that this candidate solves the unrestricted problem. Given the inductive hypothesis, the continuation value after failure is the same for every non-decisive question. Accordingly,

the current stage presents a Shannon choice problem with J_s ex ante symmetric actions, one of which yields payoff

$$1$$

and the remaining $J_s - 1$ of which yield the common payoff

$$\widehat{U}_{s-1}^I(J_s - 1).$$

By the Matějka–McKay characterization, the unrestricted optimum of this one-period Shannon problem is implemented by the symmetric one-spike posterior family described above. Hence the unrestricted Bellman problem collapses to the one-dimensional optimization

$$V_s^I(\bar{\nu}_s^I, \mathcal{Q}_s) = \max_{p \in [1/J_s, 1]} \left\{ p + (1 - p) \widehat{U}_{s-1}^I(J_s - 1) - \kappa \left[p \log(J_s p) + (1 - p) \log\left(\frac{J_s(1 - p)}{J_s - 1}\right) \right] \right\}.$$

But this is exactly the restricted sequence problem.

Step 5: Preservation of functional form. The first-order condition is

$$1 - \widehat{U}_{s-1}^I(J_s - 1) = \kappa \log\left(\frac{p(J_s - 1)}{1 - p}\right),$$

with solution

$$\hat{p}_s = \frac{e^{1/\kappa}}{e^{1/\kappa} + (J_s - 1)e^{\widehat{U}_{s-1}^I(J_s - 1)/\kappa}}.$$

Substituting back yields the recursion for $\widehat{U}_s^I(J_s)$ obtained in the restricted sequence problem. Hence

$$V_s^I(\bar{\nu}_s^I, \mathcal{Q}_s) = \widehat{U}_s^I(J_s).$$

This completes the induction. □

C.2 Sequential geometry

This appendix proves the verification result for the sequential geometry in the canonical symmetric environment.

Let $\mathcal{Q}_s$ denote the feasible question set at stage s , and let $J_s := |\mathcal{Q}_s|$. In the canonical symmetric environment, the induced diagnostic belief over the correct next question is uniform:

$$\bar{\nu}_s^S(q) = \frac{1}{J_s}, \quad q \in \mathcal{Q}_s.$$

Let $V_s^S(\nu, \mathcal{Q}_s)$ denote the value of the unrestricted Bellman problem, and let $\widehat{U}_s^S(J_s)$ denote the candidate value generated by the restricted sequence problem.

Theorem 3 (Verification for the sequential geometry). Consider the canonical symmetric sequential geometry in which, at each stage s , the induced diagnostic belief over the currently feasible questions $\mathcal{Q}_s$ is uniform, and conditional on success the continuation problem at stage $s-1$ again takes the same symmetric form.

Let $\widehat{U}_T$ denote the value generated by the restricted sequence problem, and let $V_T^S(\nu_T^S, \mathcal{Q}_T)$ denote the value of the unrestricted Bellman problem for the sequential geometry.

Then the restricted value coincides with the unrestricted value in the canonical symmetric environment:

$$V_T^S(\bar{\nu}_T^S, \mathcal{Q}_T) = \widehat{U}_T.$$

Equivalently, the symmetric one-spike posterior family that implements the restricted sequence problem attains the supremum in the unrestricted Bellman problem.

Proof. The proof proceeds by backward induction on the number of remaining inquiry opportunities.

Step 1: Final period. When $s = 0$, the agent has one final opportunity to refine beliefs and ask the correct question. If the correct question is chosen, the prize is obtained; if not, value is zero. Hence

$$V_0^S(\bar{\nu}_0^S, \mathcal{Q}_0) = \sup_{\Gamma_0 \in \mathcal{B}(\bar{\nu}_0^S)} \left[-\kappa I(\Gamma_0; \bar{\nu}_0^S) + \mathbb{E}_{\hat{\nu}_0 \sim \Gamma_0} \left(\max_{q \in \mathcal{Q}_0} \hat{\nu}_0(q) \right) \right].$$

This is again a one-period Shannon rational-inattention problem with J_0 ex ante symmetric actions. By the Matějka–McKay characterization, the unrestricted optimum is implemented by a symmetric one-spike posterior family and yields

$$V_0^S(\bar{\nu}_0^S, \mathcal{Q}_0) = \kappa \log \left(\frac{e^{1/\kappa} + J_0 - 1}{J_0} \right).$$

Equivalently, the optimal probability that the final question is correct is

$$\hat{p}_0 = \frac{e^{1/\kappa}}{e^{1/\kappa} + J_0 - 1}.$$

This is exactly the terminal value generated by the restricted sequence problem, so

$$V_0^S(\bar{\nu}_0^S, \mathcal{Q}_0) = \widehat{U}_0^S(J_0).$$

Step 2: Inductive hypothesis. Fix $s \geq 1$, and assume that at stage $s-1$ the unrestricted Bellman problem coincides with the candidate restricted value:

$$V_{s-1}^S(\bar{\nu}_{s-1}^S, \mathcal{Q}_{s-1}) = \widehat{U}_{s-1}^S(J_{s-1}).$$

Step 3: Candidate posterior family. Consider the symmetric one-spike posterior family indexed by a scalar success probability $p \in [1/J_s, 1]$. For each $q \in \mathcal{Q}_s$, define the posterior $\hat{\nu}^{q,p}$ by

$$\hat{\nu}^{q,p}(q) = p, \quad \hat{\nu}^{q,p}(q') = \frac{1-p}{J_s-1}, \quad q' \neq q.$$

Let these posteriors occur with equal ex ante probability $1/J_s$. This family is Bayes-plausible under the uniform prior $\bar{\nu}_s^S$.

In the canonical sequential geometry, if the posterior-maximizing question q is chosen and succeeds, the next-stage problem again takes the canonical symmetric form. Hence success leads to continuation value

$$\hat{U}_{s-1}^S(J_{s-1}).$$

Step 4: Guess and verify. Under the candidate posterior family, choosing the posterior-maximizing question yields current payoff

$$p \hat{U}_{s-1}^S(J_{s-1}).$$

The Shannon information cost of implementing the one-spike family is

$$\kappa \left[p \log(J_s p) + (1-p) \log\left(\frac{J_s(1-p)}{J_s-1}\right) \right].$$

Therefore the value generated by the candidate family is

$$p \hat{U}_{s-1}^S(J_{s-1}) - \kappa \left[p \log(J_s p) + (1-p) \log\left(\frac{J_s(1-p)}{J_s-1}\right) \right].$$

We now verify that this candidate solves the unrestricted problem. Given the inductive hypothesis, the continuation payoff after success is the same for every currently correct question, while choosing any incorrect question yields zero. Accordingly, the current stage presents a Shannon choice problem with J_s ex ante symmetric actions, one of which yields payoff

$$\hat{U}_{s-1}^S(J_{s-1}),$$

and the remaining $J_s - 1$ of which yield payoff

$$0.$$

By the Matějka–McKay characterization, the unrestricted optimum of this one-period Shannon problem is implemented by the symmetric one-spike posterior family described above. Hence the unrestricted Bellman problem collapses to the one-dimensional optimization

$$V_s^S(\bar{\nu}_s^S, \mathcal{Q}_s) = \max_{p \in [1/J_s, 1]} \left\{ p \hat{U}_{s-1}^S(J_{s-1}) - \kappa \left[p \log(J_s p) + (1-p) \log\left(\frac{J_s(1-p)}{J_s-1}\right) \right] \right\}.$$

But this is exactly the restricted sequence problem.

Step 5: Preservation of functional form. The first-order condition is

$$\widehat{U}_{s-1}^S(J_{s-1}) = \kappa \log\left(\frac{p(J_s - 1)}{1 - p}\right),$$

with solution

$$\hat{p}_s = \frac{e^{\widehat{U}_{s-1}^S(J_{s-1})/\kappa}}{e^{\widehat{U}_{s-1}^S(J_{s-1})/\kappa} + J_s - 1}.$$

Substituting back yields the recursion for $\widehat{U}_s^S(J_s)$ obtained in the restricted sequence problem. Hence

$$V_s^S(\bar{\nu}_s^S, \mathcal{Q}_s) = \widehat{U}_s^S(J_s).$$

This completes the induction. □

D Closed-form calculation for the identification geometry

Lemma 1 (Closed-form solution for the identification geometry). *Consider the restricted sequence problem with stages $t = 0, 1, \dots, T$, where $n_t \geq 2$ denotes the number of feasible questions at stage t , and $t = 0$ is the final stage. Let*

$$A \equiv e^{1/\kappa}.$$

Then the value function satisfies the recursion

$$\widehat{U}_t = \kappa \log\left(\frac{A + (n_t - 1)e^{\widehat{U}_{t-1}/\kappa}}{n_t}\right), \quad t \geq 1,$$

with base case

$$\widehat{U}_0 = \kappa \log\left(\frac{A + n_0 - 1}{n_0}\right),$$

and admits the closed-form expression

$$\widehat{U}_t = \kappa \log\left(A - (A - 1) \prod_{s=0}^t \frac{n_s - 1}{n_s}\right).$$

Proof. For $t = 0$, we have

$$\widehat{U}_0 = \kappa \log\left(\frac{A + n_0 - 1}{n_0}\right) = \kappa \log\left(A - (A - 1) \frac{n_0 - 1}{n_0}\right),$$

so the formula holds.

Now suppose the expression holds for $t - 1$, so that

$$e^{\widehat{U}_{t-1}/\kappa} = A - (A - 1) \prod_{s=0}^{t-1} \frac{n_s - 1}{n_s}.$$

Substituting into the recursion gives

$$\widehat{U}_t = \kappa \log \left(\frac{A + (n_t - 1) \left[A - (A - 1) \prod_{s=0}^{t-1} \frac{n_s - 1}{n_s} \right]}{n_t} \right).$$

Expanding the numerator,

$$A + (n_t - 1)A - (n_t - 1)(A - 1) \prod_{s=0}^{t-1} \frac{n_s - 1}{n_s} = n_t A - (n_t - 1)(A - 1) \prod_{s=0}^{t-1} \frac{n_s - 1}{n_s}.$$

Dividing by n_t yields

$$A - (A - 1) \frac{n_t - 1}{n_t} \prod_{s=0}^{t-1} \frac{n_s - 1}{n_s} = A - (A - 1) \prod_{s=0}^t \frac{n_s - 1}{n_s}.$$

Hence

$$\widehat{U}_t = \kappa \log \left(A - (A - 1) \prod_{s=0}^t \frac{n_s - 1}{n_s} \right),$$

which completes the induction. □

E A Posterior Calculation for the Binary Shannon Problem

The following calculation clarifies the appearance of the Shannon log-sum-exp operator in the main text. In the symmetric binary case, the posterior representation of the Shannon problem makes the cancellation between expected utility and optimized information cost particularly transparent.

Setup. Let the payoff-relevant state be

$$s \in \{0, 1\},$$

with prior belief

$$\mu = \left(\frac{1}{2}, \frac{1}{2} \right).$$

There are two actions, A and B . Action A is correct in state 1, and action B is correct in state 0, so payoffs are

$$u(A, 1) = 1, \quad u(B, 0) = 1, \quad u(A, 0) = 0, \quad u(B, 1) = 0.$$

The decision maker chooses an information structure at Shannon cost $\kappa I(s; z)$ and then selects the action that maximizes expected utility given the posterior belief.

Posterior representation. With a binary state and two actions, any symmetric information structure (such as the optimal structure) can be represented by two posterior beliefs

$$\gamma^A = (\gamma, 1 - \gamma), \quad \gamma^B = (1 - \gamma, \gamma),$$

corresponding to signals recommending actions A and B respectively.

Bayes plausibility requires the prior to equal the average posterior:

$$\frac{1}{2}\gamma^A + \frac{1}{2}\gamma^B = \left(\frac{1}{2}, \frac{1}{2}\right).$$

Thus the two posteriors occur with equal probability.

Expected utility. After observing γ^A , the agent chooses action A , which is correct with probability γ . After observing γ^B , the agent chooses action B , which is also correct with probability γ . Hence expected utility is

$$EU = \gamma.$$

Shannon information cost. Mutual information can be written as prior entropy minus expected posterior entropy:

$$I(s; z) = H(\mu) - \mathbb{E}[H(\gamma)].$$

Since the prior is $(1/2, 1/2)$,

$$H(\mu) = \log 2.$$

Both posteriors have the same entropy

$$H(\gamma) = -\gamma \log \gamma - (1 - \gamma) \log(1 - \gamma),$$

so

$$I(s; z) = \log 2 + \gamma \log \gamma + (1 - \gamma) \log(1 - \gamma).$$

The Shannon information cost is therefore

$$\kappa I = \kappa \log 2 + \kappa \gamma \log \gamma + \kappa(1 - \gamma) \log(1 - \gamma).$$

Optimal posterior. Under Shannon costs the interior optimal posterior takes the form

$$\gamma = \frac{e^{1/\kappa}}{1 + e^{1/\kappa}}.$$

Let

$$Z = 1 + e^{1/\kappa}.$$

Then

$$\gamma = \frac{e^{1/\kappa}}{Z}, \quad 1 - \gamma = \frac{1}{Z}.$$

Hence

$$\log \gamma = \frac{1}{\kappa} - \log Z, \quad \log(1 - \gamma) = -\log Z.$$

Substituting into the entropy expression gives

$$\gamma \log \gamma + (1 - \gamma) \log(1 - \gamma) = \frac{\gamma}{\kappa} - \log Z.$$

Thus

$$I = \log 2 + \frac{\gamma}{\kappa} - \log Z,$$

and multiplying by κ yields

$$\kappa I = \kappa \log 2 + \gamma - \kappa \log Z.$$

Optimized value. Subtracting the information cost from expected utility gives

$$EU - \kappa I = \gamma - (\kappa \log 2 + \gamma - \kappa \log Z) = \kappa \log Z - \kappa \log 2.$$

Since $Z = 1 + e^{1/\kappa}$,

$$V = \kappa \log \left(\frac{1 + e^{1/\kappa}}{2} \right).$$

This is the symmetric ($\pi = \frac{1}{2}$) case of the Shannon log-sum-exp expression

$$V(\pi) = \kappa \log \left((1 - \pi) + \pi e^{1/\kappa} \right).$$

The calculation illustrates how the posterior formulation produces the familiar log-sum-exp

operator once expected utility and the optimized Shannon information cost are combined. In the main text, the same operator appears because the dynamic inquiry problem reduces to a binary success margin—whether the decisive question is located— and Shannon refinement transforms the baseline probability of success through this operator.

Remark. The same calculation extends to arbitrary priors $\mu = (\pi, 1 - \pi)$. In that case the Shannon identity yields

$$V(\pi) = \kappa \log\left((1 - \pi) + \pi e^{1/\kappa}\right),$$

so the symmetric calculation above corresponds to the special case $\pi = \frac{1}{2}$. The symmetric case is presented because it makes the posterior geometry particularly transparent.