Planning under AI:
Modeling and Measuring Question-Driven Learning

Andrew Caplin*

March 10, 2026

Abstract

Economic models of information were developed for environments in which reliable information about the world was the key constraint. Search theory studies the cost of acquiring such information, and related frameworks analyze how agents process signals about payoff-relevant states before acting. In these models the informational structure of the environment is known: agents learn about the state of the world but not about how information should be generated or interpreted.

This paper studies decision-making when the informational structure of inquiry itself is uncertain. Agents choose questions before committing to an action, but the diagnostic value of those questions is not known in advance. Questions generate answers through an uncertain diagnostic structure that determines which inquiries are informative and how answers should be interpreted.

The model treats questions as primitive acts of inquiry and characterizes planning as a dynamic problem of organizing investigation under diagnostic uncertainty. Agents update beliefs about both payoff states and diagnostic regimes while allocating cognitive effort to evaluating which questions are worth asking.

Artificial intelligence makes this margin economically consequential and empirically observable. When candidate answers can be generated at negligible cost, the scarcity that constrains decision-making shifts from the acquisition of information to the disciplined organization of inquiry. The framework developed here provides a formal structure for analyzing how agents organize investigation in environments where answers are abundant but understanding remains scarce.

1 Introduction

Economic decisions are rarely made in a single step. They are preceded by inquiry.

*Highly Preliminary. This paper was developed during an intensive period of structured planning, much of it conducted through interaction with AI systems. The resulting externalized inquiry—question sequencing, revision, and counterfactual exploration—helped clarify the mechanism analyzed here. I am grateful to Søren Leth-Petersen and Ruth Wyatt for helpful discussions and comments.

Workers exploring career transitions investigate industries, training requirements, and geographic opportunities. Firms evaluating strategic investments examine markets, technologies, and competitors before committing resources. Researchers confronting new empirical phenomena explore alternative explanations and potential tests. In each of these environments, the quality of the final decision depends not only on the information available to the decision-maker but also on how the investigation preceding action is organized.

Which questions are asked, in what order they are asked, and how answers reshape subsequent inquiry often determine whether the investigation converges on the decisive issue or drifts across the surface of the environment.

Historically, the central constraint in this process was the cost of obtaining reliable information. Workers searched for job opportunities by contacting firms or consulting professional networks. Consumers compared prices across stores. Firms gathered market intelligence through reports, consultants, and experimentation. Each additional answer required time, effort, and often substantial expense.

Classical information economics reflects this constraint. Search models study the cost of obtaining information about prices or opportunities, bandit models analyze experimentation across known alternatives, and rational inattention models attach cognitive cost to processing signals about payoff-relevant states. In these frameworks the informational structure of the environment is taken as given. Agents learn about the state of the world but not about how information should be generated or interpreted.

In many real environments this assumption is unrealistic. Decision makers often face uncertainty not only about the state of the world but also about which questions are informative. A worker exploring a new career path may not know which aspects of the industry determine entry. A firm evaluating an investment may not know which market indicators reveal the true prospects of a technology. Researchers confronting new data may not know which empirical tests will isolate the decisive mechanism.

Planning therefore involves organizing inquiry when the diagnostic value of available questions is itself uncertain.

This paper develops a framework for analyzing decision-making under such *diagnostic uncertainty*. The model treats questions as primitive acts of inquiry. Each question generates answers through a diagnostic structure that determines how answers relate to the underlying payoff state. When the diagnostic structure is unknown, agents must learn not only about the payoff-relevant state of the world but also about which questions are informative and how answers should be interpreted.

Planning becomes a dynamic problem of organizing investigation. Agents select questions sequentially before committing to an action, update beliefs about both payoff states and diagnostic regimes, and allocate cognitive effort to evaluating which lines of inquiry are worth pursuing. The framework developed here—*Question-Driven Learning*—characterizes how inquiry unfolds when the informational value of questions is itself uncertain.

Artificial intelligence alters the economic importance of this margin. Modern AI systems generate candidate answers to an extraordinary range of questions almost instantly. They summarize documents, propose explanations, compare alternatives, and suggest new directions of investigation. Information that once required hours of search or consultation can now be produced within seconds.

When answers become inexpensive to generate, the scarcity that constrains decision-making shifts. The bottleneck is no longer the acquisition of information but the cognitive effort required to organize inquiry. The problem moves from scarce information to scarce cognition.

Artificial intelligence also makes this margin empirically observable. Interactive AI systems record the sequence of prompts and responses that precedes a decision. These interaction traces preserve the structure of inquiry: the order of questions, the branching of investigation, and the point at which inquiry converges.

The framework developed in this paper provides a structure for interpreting such traces. Inquiry sequences can be understood as realizations of a dynamic process in which agents select questions under uncertainty about diagnostic structure. Planning under diagnostic uncertainty therefore becomes measurable.

The remainder of the paper develops these ideas formally. Section 2 introduces the architecture of planning under diagnostic uncertainty. Questions are modeled as acts of inquiry, diagnostic structures determine how questions generate answers, and agents update beliefs about both payoff states and diagnostic regimes. Section 3 introduces diagnostic refinement, which captures costly reasoning about which questions are informative. Section 4 embeds these elements in a dynamic program describing how inquiry unfolds prior to commitment.

Section 6 analyzes two canonical diagnostic geometries that isolate distinct planning margins. Section 7 discusses how planning becomes measurable through inquiry traces and structured decision environments. Section 8 examines the role of planning in decision-making in environments shaped by artificial intelligence.

Together these sections establish planning under diagnostic uncertainty as a formal domain of economic analysis.

A fuller discussion of related intellectual traditions appears in Section 10.

2 The Architecture of Enquiry

Consequential decisions require individuals to gather information before acting. In classical models of information acquisition the informational properties of available signals are known. Agents may choose which signals to observe, but they do not need to learn how those signals should be interpreted.

In many real environments this assumption is unrealistic. Decision makers often face uncertainty not only about the state of the world but also about which questions are informative. Planning therefore involves organizing inquiry when the diagnostic value of available questions is itself un-

certain.

The model developed in this paper provides a framework for analyzing such environments. Questions generate answers through an unknown diagnostic structure, and agents may devote cognitive effort to refining beliefs about which questions reveal useful information.

Structure of the model. The model proceeds in three steps. Section 2 introduces the primitives of the environment and the role of diagnostic structure in determining how questions generate answers. Section 3 introduces diagnostic refinement, which captures costly reasoning about which questions are informative. Section 4 embeds these objects in a dynamic program describing how inquiry unfolds before commitment.

2.1 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

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

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

2.2 Questions and Answers

Before committing to an action, the agent may acquire information through inquiry.

In practice, inquiry may take many forms. An individual may reflect internally, consult an advisor, examine a dataset, or pose a prompt to a computational system. Such processes may be conscious or subconscious, internal or external. The model abstracts from these details by representing inquiry through discrete acts that generate observable informational outcomes.

Assumption 1 (Admissible Questions)**.** *Inquiry proceeds through a finite set of admissible questions*

$$\mathcal{Q} = \{q_1, \ldots, q_K\}.$$

Each question $q \in \mathcal{Q}$ produces an observable answer drawn from a finite answer set

$$a \in A_q.$$

Questions are primitive objects in the model. They represent admissible acts of inquiry available to the agent.

A question is characterized by two elements:

- the distinction it investigates, and
- the source from which the answer is obtained.

Thus the same verbal query directed to different sources constitutes different questions. Consulting Advisor A rather than Advisor B, querying Dataset 1 rather than Dataset 2, or posing the same prompt to different language models correspond to distinct elements of $\mathcal{Q}$.

This reflects a basic feature of planning: the informational value of a question depends not only on what is asked but on where the answer comes from.

The set $\mathcal{Q}$ reflects the feasibility of articulating and executing questions. In practice, inquiry is constrained by linguistic, institutional, and technological factors. Individuals cannot construct arbitrary informational experiments; they pose questions that can realistically be formulated, communicated, and answered using available tools or sources of expertise.

2.3 Diagnostic Structure

In many environments the informational value of a question is itself uncertain.

To represent this formally we introduce a diagnostic structure

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

The diagnostic structure governs how answers should be interpreted. It determines how answers generated by each question relate to the underlying payoff state.

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

Subjective representations. The framework is entirely subjective in the Savage (1954) tradition. 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.

This subjectivity allows the model to accommodate rich individual heterogeneity in inquiry. Two individuals facing the same environment may ask different questions, consult different sources, or interpret the same answers through different diagnostic lenses. Later sections exploit this feature to analyze differences in planning protocols and inquiry performance across agents.

2.4 Answer Generation as Experiments

Conditional on a diagnostic structure D, each question induces a standard Blackwell experiment.

For each question q the distribution of answers is given by

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

Equivalently, fixing D defines an experiment mapping

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

If the diagnostic structure were known, every question would correspond to a well-defined statistical experiment.

When the diagnostic structure is uncertain, however, a question does not correspond to a single experiment. Instead it represents a *lottery over experiments* indexed by D. Planning therefore requires reasoning not only about the payoff state of the world but also about how answers should be interpreted.

2.5 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

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

More generally, let

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

denote a prior belief at some stage of inquiry.

After asking question q and observing answer $a \in A_q$, beliefs update to a posterior. To emphasize timing, we denote posteriors by γ rather than by μ.

Bayesian update operator. For prior belief μ, question q, and answer $a \in A_q$, define the posterior belief

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

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 may therefore provide information about both the payoff state θ and the diagnostic structure D governing the interpretation of answers.

2.6 Indivisibility and Feasible Inquiry

Questions are treated as indivisible acts of inquiry. Once posed, a question generates an answer from its associated answer set A_q.

A single question may investigate multiple informational distinctions simultaneously. For example, a statistical test may bear on several hypotheses, a consultation with an advisor may address several possible explanations, and a single prompt to a language model may screen multiple competing interpretations.

What constrains inquiry is therefore not the logical separability of ideas, but the feasibility of formulating and executing questions. The set $\mathcal{Q}$ represents the menu of questions that can realistically be posed. This menu reflects linguistic, institutional, and technological constraints on inquiry. Agents cannot freely construct arbitrary experiments; they choose among available advisors, available datasets, available instruments, and available analytic procedures.

Thus $\mathcal{Q}$ captures the feasible space of inquiry acts available to the agent.

2.7 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. After the inquiry phase concludes, the agent commits to an action

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

The horizon T captures the limited time, attention, and cognitive resources available for inquiry.

Key implication. When diagnostic structure is uncertain, the object of planning is not only what to learn about θ, but also how to organize inquiry when the diagnostic value of questions is itself uncertain. The next section introduces a formal representation of diagnostic planning.

3 Diagnostic Planning

The preceding section introduced the basic elements of 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.

These primitives describe how inquiry generates information. They do not yet describe how agents decide *which questions are worth asking.*

When the diagnostic value of questions is uncertain, planning requires reasoning about the informational structure of inquiry itself. Agents must evaluate which distinctions are likely to reveal useful information and how answers may reshape subsequent investigation.

This section introduces a representation of that process.

The key modeling move is to attach cognitive cost not to learning about the payoff state, but to refining beliefs about diagnostic structure. Planning therefore involves allocating cognitive effort

to evaluating which questions are informative before selecting the next inquiry.

3.1 A Shift in the Domain of Costly Cognition

Economic models of information have traditionally attached cost to acquiring or processing information about payoff-relevant states.

Search models study the cost of obtaining information about prices or opportunities. Bandit models analyze the allocation of sampling across known experiments. Rational inattention models attach cost to the precision with which signals about states of the world are processed.

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

The present framework shifts the domain of costly cognition.

Rather than attaching cost to learning about the payoff state θ, the model attaches cost to evaluating *diagnostic structure*: determining which questions are informative and how sequences of questions reshape the decision problem.

In this sense the model moves the object of costly cognition one level higher in the hierarchy of decision-making.

Search theory studies the cost of obtaining information. Rational inattention studies the cost of processing information. Question-Driven Learning studies the cost of organizing inquiry.

3.2 Planning as Anticipatory Reasoning

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

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

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

Psychologist George Kelly famously described individuals as *anticipatory systems* that construct internal representations of the world in order to predict how future events may unfold.

The model developed here does not attempt to represent this internal reasoning process explicitly. Instead it provides a formal representation of the outcome of that reasoning.

3.3 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 corresponds to 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 beliefs after observing answers is treated as routine. The key margin concerns costly refinement of beliefs about diagnostic structure.

3.4 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 distribution over posterior beliefs about diagnostic structure and ultimately choice of action.

Define

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

The set $\mathcal{B}(\nu_t)$ contains all Bayes-plausible distributions of posterior beliefs about diagnostic structure.

An element

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

represents a lottery over posterior diagnostic beliefs. Drawing

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

produces the refined belief

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

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

3.5 Cost of Diagnostic Refinement

Refining beliefs about diagnostic structure requires cognitive effort.

We represent this effort by a cost functional

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

The degenerate refinement that leaves beliefs unchanged incurs no cost:

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

The cost function $\mathcal{C}$ therefore captures the cognitive resources required to evaluate which questions are informative.

Different specifications of $\mathcal{C}$ correspond to different assumptions about the nature of cognitive effort. Later sections consider structured cases such as time-separable costs and Shannon information costs.

3.6 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. Refinement therefore concentrates belief updates on those diagnostic distinctions most likely to influence the optimal choice of the next question.

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 informational structure.

In this sense refinement determines how precisely agents learn which questions are worth asking.

3.7 Planning Skill

Planning skill arises from the ability to allocate diagnostic attention efficiently.

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.

Several forms of planning failure may therefore arise:

- **Shallow refinement**: agents fail to recognize which questions are likely to be informative.
- **Mis-sequencing**: agents pursue diagnostically weak questions while overlooking more informative sources.
- **Wasted anticipation**: agents expend cognitive effort analyzing contingencies that do not affect the next inquiry step.

- **Premature stopping**: agents abandon inquiry before identifying useful sources of information.
- **Diffuse exploration**: Diffuse exploration: attention is spread across too many sources rather than focused 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 their beliefs about which questions are informative.

Planning skill therefore reflects not superior Bayesian updating about payoff states, but the ability to learn which questions are worth asking.

3.8 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 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 rational-inattention problem: the object of attention is not only the state of the world but the informational structure governing how questions generate answers.

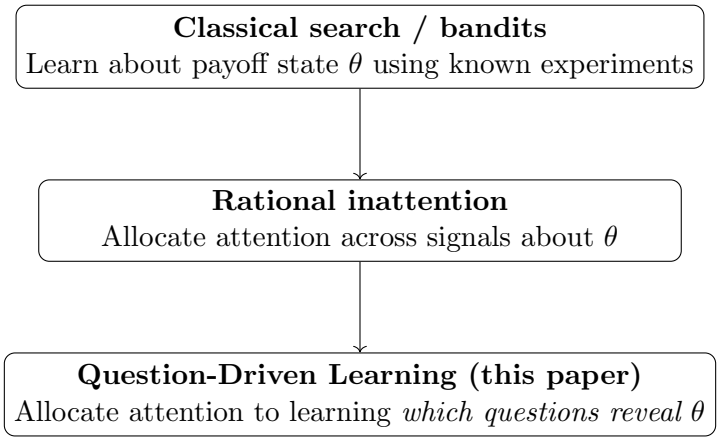


Figure 1: A shift in the object of attention. Classical models treat the experiment mapping as known. Question-Driven Learning endogenizes learning about diagnostic structure.

3.9 From Diagnostic Planning to Sequential Inquiry

The preceding discussion introduces 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.

This refinement process determines how precisely agents understand the informational value of alternative questions. When refinement costs are low, beliefs about diagnostic structure can be tailored closely to the environment. When refinement costs are high, beliefs remain coarse and question choice becomes less responsive to the true informational structure.

Planning therefore involves the sequential allocation of cognitive effort across inquiry steps.

At each stage the agent faces two decisions: 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 on both the information it may reveal about the payoff state and the possibility that its answer may reshape the structure of subsequent inquiry.

The next section formalizes this sequential planning problem as a dynamic program.

4 Dynamic Structure of Inquiry

The preceding sections introduced the primitives of the environment and the diagnostic refinement operator. We now embed these objects in a finite-horizon dynamic program describing how inquiry unfolds before commitment.

4.1 Timing and State

There is a finite horizon T. Time is indexed by periods remaining

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

At the beginning of period t the agent holds a prior belief

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

which summarizes all information generated by past inquiry.

Let

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

denote the marginal belief over diagnostic structures.

The agent also faces a set of remaining admissible questions

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

4.2 Within-Period Belief Evolution

Beliefs evolve within a period through three stages:

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

- μ_t is the prior belief entering period t.
- Diagnostic refinement produces a refined belief

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

  where $\hat{\nu}_t$ is drawn from the refinement distribution Γ_t.
- After question q is asked and answer a is observed, beliefs update to

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

Thus $_t$ summarizes the informational consequences of past inquiry, γ_t reflects the refined diagnostic belief used to evaluate the next question, and μ_{t-1} becomes the prior belief entering the next period.

4.3 Decision Structure

At each period $t \geq 1$, the agent faces three nested decisions:

1. **Stop and commit** to an action

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

2. **Refine diagnostic beliefs** by selecting

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

3. **Choose the next question**

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

   A refinement draw

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

   generates the refined belief γ_t. Conditional on γ_t, the agent selects question q.

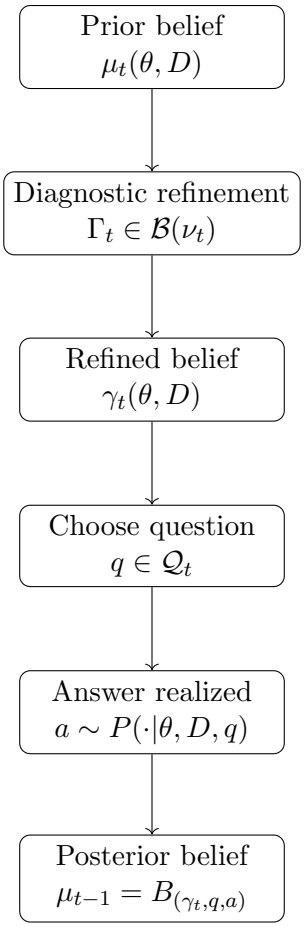


Figure 2: Architecture of inquiry within a period. The agent enters the period with prior belief μ_t, may refine beliefs about diagnostic structure, selects a question, observes an answer, and updates beliefs for the next period.

4.4 Answer Generation and Updating

After question q is asked, an answer

$$a \in A_q$$

is observed.

Posterior beliefs are obtained using the Bayesian update operator

$$B_{(\gamma,q,a)}.$$

Formally,

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

The resulting belief becomes the prior for the next period:

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

4.5 Continuation Operator

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

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

where

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

4.6 Bellman Recursion

The stopping value is

$$V_0(\mu, \mathcal{Q}) = \max_{y \in \mathcal{Y}} \sum_{\theta} \mu^{\Theta}(\theta) U(y, \theta),$$

where μ^{Θ} denotes the marginal belief over payoff states. Since payoffs depend only on the payoff state, only the marginal belief over θ enters the stopping value.

Define the continuation value

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

The value function satisfies

$$V_t(\mu_t, \mathcal{Q}_t) = \max \left\{ V_0(\mu_t, \mathcal{Q}_t), J_t(\mu_t, \mathcal{Q}_t) \right\}.$$

4.7 SDSC Representation and Existence

The dynamic program above is written in terms of posterior beliefs generated by diagnostic refinement. At time t the agent holds prior belief μ_t over (θ, D), refines diagnostic beliefs through a refinement Γ_t, obtains a refined belief γ_t, chooses a question q, and after observing answer a updates beliefs to

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

Because the choice of question ultimately depends on the realized diagnostic belief, it is convenient to represent the refinement stage directly in terms of the stochastic rule that determines which question is asked for each diagnostic structure as in Matejka and McKay (2015). This leads to a representation in terms of state-dependent stochastic choice (SDSC).

Diagnostic-state–dependent question choice. Let

$$\sigma : \mathcal{D} \to \Delta(\mathcal{Q}_t)$$

denote a stochastic choice rule that assigns to each diagnostic structure D a probability distribution over admissible questions. For each D, the rule $\sigma(\cdot|D)$ specifies the probability that question $q \in \mathcal{Q}_t$ is selected.

Given belief ν_t over diagnostic structures, such a rule generates the expected continuation value

$$\sum_{D \in \mathcal{D}} \nu_t(D) \sum_{q \in \mathcal{Q}_t} \sigma(q|D)\, W_{t-1}(\gamma_t, q, \mathcal{Q}_t).$$

Not every SDSC rule can be implemented without diagnostic refinement. To condition question choice on D, the agent must acquire information about diagnostic structure. We therefore characterize the minimal cost of implementing a given SDSC rule.

Axioms on refinement costs. Let $\mathcal{C}(\nu, \Gamma)$ denote the cost of implementing refinement $\Gamma \in \mathcal{B}(\nu)$. We impose two mild conditions.

Assumption 2 (Information monotonicity of refinement costs)**.** *For each diagnostic belief ν and refinement $\Gamma \in \mathcal{B}(\nu)$:*

1. ***Zero cost of inattention.*** *The degenerate refinement that leaves beliefs unchanged is costless:*
$$\mathcal{C}(\nu, \Gamma^{\circ}_{\nu}) = 0.$$
2. ***Monotonicity in the Blackwell order.*** *If refinement Γ_1 is Blackwell more informative than Γ_2, then*
$$\mathcal{C}(\nu, \Gamma_1) \geq \mathcal{C}(\nu, \Gamma_2).$$

These conditions impose the standard discipline used in information-cost models. More informative refinements require weakly greater cognitive effort, while inattention is costless.

Induced SDSC cost. For any SDSC rule σ and belief ν, define the induced cost

$$C_t(\sigma; \nu) = \inf_{\Gamma \in \mathcal{B}(\nu)} \{\mathcal{C}(\nu, \Gamma) : \Gamma \text{ induces } \sigma\}.$$

Thus $C_t(\sigma; \nu)$ is the minimum cost of implementing rule σ through some diagnostic refinement.

Regularity. We impose a standard regularity condition on the induced SDSC cost.

Assumption 3 (Regularity of the induced SDSC cost)**.** *For each belief ν and rule σ:*

1. *the infimum defining $C_t(\sigma; \nu)$ is attained by some refinement $\Gamma^{\star}(\sigma, \nu)$;*

2. the mapping $(\sigma, \nu) \mapsto C_t(\sigma; \nu)$ is continuous.

These conditions are satisfied for a large class of cost functions used in rational-inattention models. In particular they hold for all posterior-separable costs of the form

$$\mathcal{C}(\nu, \Gamma) = \mathbb{E}_{\hat{\nu} \sim \Gamma}[\phi(\hat{\nu})] - \phi(\nu),$$

where ϕ is continuous and concave on the simplex $\Delta(\mathcal{D})$. Shannon mutual-information costs arise as a canonical member of this class as in Caplin et al. (2022).

SDSC formulation of the continuation problem. Under these definitions the continuation value in the dynamic program can be written as

$$J_t(\mu_t, \mathcal{Q}_t) = \sup_{\sigma} \left[\sum_{D \in \mathcal{D}} \nu_t(D) \sum_{q \in \mathcal{Q}_t} \sigma(q|D)\, W_{t-1}(\gamma_t, q, \mathcal{Q}_t) - C_t(\sigma; \nu_t) \right].$$

The refinement stage can therefore be interpreted as choosing a diagnostic-state–dependent question rule subject to the induced cost $C_t(\sigma; \nu_t)$.

Existence of an optimal policy.

Theorem 1 (Existence)**.** *Suppose Θ, $\mathcal{D}$, $\mathcal{Q}$, and $\mathcal{Y}$ are finite, utility is bounded, and Assumptions 2 and 3 hold. Then for every state $(\mu_t, \mathcal{Q}_t)$ the value function $V_t(\mu_t, \mathcal{Q}_t)$ is well defined and an optimal policy exists.*

Proof. For fixed $(\mu_t, \mathcal{Q}_t)$ the SDSC choice set

$$\{\sigma : \mathcal{D} \to \Delta(\mathcal{Q}_t)\}$$

is compact because both $\mathcal{D}$ and $\mathcal{Q}_t$ are finite. By Assumption 3, the induced cost $C_t(\sigma; \nu_t)$ is continuous. The continuation payoff is linear in σ. Hence the objective function is continuous on a compact domain and therefore attains its maximum. Backward induction over the finite horizon establishes existence. □

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.

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 the theoretical 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 does not doubt the existence of a true demand state, but 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 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. In each case the decision maker must determine which questions are likely to generate useful information when the diagnostic structure of inquiry is uncertain.

The general model developed earlier allows arbitrary diagnostic mappings

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

and arbitrary diagnostic refinement strategies. In such environments inquiry can take many forms.

To clarify the economic forces governing optimal planning, it is useful to analyze simplified environments that isolate distinct planning margins. The following section introduces two canonical diagnostic geometries that capture different ways in which diagnostic uncertainty can shape inquiry.

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 environments described above illustrate that diagnostic planning can take different forms depending on how questions reveal information.

The Bellman recursion in Section 4 allows arbitrary diagnostic mappings

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

and arbitrary refinement strategies. To clarify the economic forces governing optimal inquiry, we analyze two canonical diagnostic geometries that isolate distinct margins of inquiry.

In *identification environments*, exactly one question resolves the payoff-relevant uncertainty, so the planning problem is to identify which question is decisive. In *sequential-discipline environments*, value arises from maintaining a correct sequence of inquiry, so the planning problem is to select the correct next question at each stage. Both environments share the same dynamic program and cost structure of Section 4; only the informational geometry governing how questions reveal diagnostic structure differs. In both geometries the agent's problem reduces to choosing the probability of success in the next step of the diagnostic process.

6.1 Identification Geometry: One Decisive Question

The first canonical geometry captures environments in which exactly one question resolves the payoff-relevant uncertainty, but the agent does not know which question it is. The planning problem is therefore to identify which question is decisive.

This structure arises naturally in many settings. A manager may have several analyses available but only one resolves the key uncertainty. An entrepreneur may consult several advisors but only one possesses the relevant expertise. In such environments inquiry fails in two ways: by asking the wrong question, or by allocating attention to distinctions among questions that will not be asked.

Environment

Assumption 4 (Unique decisive question)**.** *For each diagnostic structure $D \in \mathcal{D}$ there exists a unique question $r(D) \in \mathcal{Q}$ such that:*

- *If $q = r(D)$, observing its answer reveals the payoff-relevant state θ completely.*
- *If $q \neq r(D)$, the distribution of answers is independent of θ.*

Let

$$R = r(D)$$

denote the identity of the decisive question.

Assumption 5 (Uniform prior over decisive questions)**.** *Conditional on the current candidate set $\mathcal{Q}_t$ with $J_t = |\mathcal{Q}_t|$, the decisive question is equally likely to be any remaining candidate:*

$$\Pr(R = q) = \frac{1}{J_t}.$$

Assumption 6 (Shannon refinement costs)**.** *Refining beliefs over candidate questions incurs Shannon (mutual-information) costs. If posterior beliefs over the J candidate questions are $q \in \Delta_J$ and the prior is π, the cost of refinement is*

$$\kappa D(q\|\pi) = \kappa \sum_{i=1}^{J} q_i \log\left(\frac{q_i}{\pi_i}\right),$$

where $\kappa > 0$ measures the marginal cost of attention.

Under Assumption 4, asking a non-decisive question simply eliminates that candidate without providing information about the remaining questions.

Bellman problem

Let $V_t(J)$ denote the maximal expected net payoff when J candidate questions remain and t periods remain.

If the decisive question is asked the agent learns the payoff-relevant state and obtains payoff normalized to 1. If a non-decisive question is asked the candidate is eliminated and the problem continues with $J-1$ remaining candidates.

Let

$$x(q) = \max_i q_i$$

denote the probability that the question selected after learning is decisive.

Under Assumptions 5–6, the Bellman problem is

$$V_t(J) = \max_{q \in \Delta_J} \Big[x(q) + (1 - x(q)) V_{t-1}(J-1) - \kappa D(q \| \pi) \Big].$$

The stopping value is

$$V_0 = \max_{y \in \mathcal{Y}} \mathbb{E}[U(y, \theta)].$$

Theorem 2 (Optimal refinement under identification geometry)**.** *Under Assumptions 4–6, there exists an optimal refinement policy that produces a posterior of the form*

$$\left(x, \frac{1-x}{J-1}, \ldots, \frac{1-x}{J-1} \right)$$

for some success probability x.

The optimal success probability satisfies

$$x_t^*(J) = \frac{1}{1 + (J-1) \exp\left(-\frac{1 - V_{t-1}(J-1)}{\kappa} \right)}.$$

The associated value recursion is

$$V_t(J) = \kappa \log \left(\frac{e^{1/\kappa} + (J-1) e^{V_{t-1}(J-1)/\kappa}}{J} \right).$$

Proof strategy The proof proceeds in two steps. First we solve a restricted learning problem in which refinement policies produce a single candidate question. We then show that for any success probability x, this restricted posterior uniquely minimizes Shannon information costs. Consequently the restricted problem yields the global optimum of the original Bellman problem.

A restricted learning problem

The key step is to show that for any success probability x, the posterior considered below uniquely minimizes Shannon information costs. We therefore begin by solving a restricted learning problem

To solve the Bellman problem we first analyze a restricted class of refinement policies in which attention is used only to identify a single best candidate question.

Under such a policy the posterior takes the form

$$q(x) = \left(x, \frac{1-x}{J-1}, \ldots, \frac{1-x}{J-1}\right), \qquad x \in [1/J, 1].$$

Economically, this means attention is allocated solely toward distinguishing one promising candidate question from the remaining symmetric alternatives. No attention is spent refining distinctions among questions that will not be asked in the current period.

We first solve the Bellman problem under this restricted class of learning policies. We then show that any refinement policy generating the same success probability x must incur strictly higher information cost. Consequently the restricted problem yields the global solution to the original Bellman problem.

Restricted optimization

Under the posterior $q(x)$ the Shannon cost becomes

$$C_J(x) = \kappa \left[x \log(Jx) + (1-x) \log\left(\frac{J(1-x)}{J-1}\right)\right].$$

The restricted Bellman problem is therefore

$$V_t(J) = \max_{x \in [1/J, 1]} \Big\{ x + (1-x) V_{t-1}(J-1) - C_J(x) \Big\}.$$

Proposition 1 (Optimal learning intensity)**.** *The optimal success probability satisfies*

$$1 - V_{t-1}(J-1) = \kappa \log\left(\frac{x(J-1)}{1-x}\right).$$

Equivalently,

$$x_t^*(J) = \frac{1}{1 + (J-1)\exp\Big(-\frac{1 - V_{t-1}(J-1)}{\kappa}\Big)}.$$

Proposition 2 (Value recursion)**.** *The value function satisfies*

$$V_t(J) = \kappa \log\left(\frac{e^{1/\kappa} + (J-1)e^{V_{t-1}(J-1)/\kappa}}{J}\right), \qquad V_0 = \max_{y \in \mathcal{Y}} \mathbb{E}[U(y, \theta)].$$

Optimality of the restricted policy

The restricted problem above characterizes the optimal success probability x. We now show that the restricted posterior is the cost-minimizing way to generate any given success probability.

Lemma 1 (Cost-minimizing posterior)**.** *Fix $x \in [1/J, 1]$. Among all posterior distributions assigning probability x to the chosen question, the unique minimizer of $D(q\|\pi)$ assigns equal probability*

to the remaining $J-1$ questions:

$$q^*(x) = \left(x, \frac{1-x}{J-1}, \ldots, \frac{1-x}{J-1}\right).$$

Proof. For fixed x, minimizing relative entropy is equivalent to minimizing

$$\sum_{i \neq j} q_i \log(J q_i)$$

subject to $\sum_{i \neq j} q_i = 1 - x$. Because the function $u \mapsto u \log(Ju)$ is strictly convex, Jensen's inequality implies that the minimum occurs when all q_i are equal. □

Theorem 3 (Global optimality)**.** *Under Assumptions 4–6, the restricted refinement policy characterized above is globally optimal in the original Bellman problem.*

Proof. Any refinement policy induces a success probability x. For that same x, Lemma 1 shows that the posterior $q^*(x)$ uniquely minimizes information cost. Since the restricted problem optimizes directly over x, it therefore achieves the maximal value attainable in the original Bellman problem. □

Economic interpretation

Identification geometry isolates a pure recognition margin. Inquiry fails either by asking the wrong question or by misallocating attention across questions that will not be asked. Because all non-decisive questions have identical consequences for future inquiry, distinctions among them have no value.

Optimal inquiry therefore concentrates attention exclusively on distinguishing one candidate question from the remaining symmetric alternatives. The optimal degree of learning depends on three forces:

- the cost of refinement κ,
- the number of remaining candidates J,
- the continuation value $V_{t-1}(J-1)$.

Lower refinement costs increase the probability that the decisive question is identified immediately. When many candidates remain, beliefs are less concentrated because the prior probability that any one question is decisive is small. When continuation value is high, the agent moderates current learning effort because future opportunities to identify the decisive question remain valuable.

The identification geometry isolates a recognition margin: inquiry succeeds once the decisive question is identified. We now turn to environments in which value arises not from identifying a decisive question but from maintaining the correct trajectory of inquiry.

6.2 Sequential Discipline Geometry: One Effective Path

The second canonical geometry captures environments in which value arises from maintaining a correct sequence of inquiry. Rather than a single decisive question, there exists a unique sequence of questions that leads to resolution of the payoff-relevant uncertainty.

In such environments the challenge is not identifying a single decisive question but maintaining *discipline in inquiry*. At each stage there is exactly one next question that preserves the effective path.

This structure arises naturally in many planning problems. A medical diagnosis may require a specific sequence of tests. A scientific investigation may require pursuing a correct line of reasoning through several intermediate steps. Deviating from the correct sequence does not resolve the underlying uncertainty and may lead the investigation into an unproductive direction.

Inquiry therefore faces two failure modes. The first is asking the wrong next question, thereby stepping off the effective path. The second is misallocating attention across candidate next questions, expending cognitive effort distinguishing among questions that all lead away from the effective path.

Environment

Assumption 7 (Unique effective path)**.** *For each diagnostic structure $D \in \mathcal{D}$ there exists a unique sequence of questions*

$$(q_1^*(D), q_2^*(D), \ldots, q_T^*(D))$$

such that following this sequence reveals the payoff-relevant state θ.

At stage k the correct next question is $q_k^*(D)$. Asking this question preserves the effective path. Any other question moves the agent to an off-path continuation that does not resolve the payoff uncertainty. All off-path continuations yield the same continuation value.

Let J_k denote the number of feasible next questions at stage k.

Assumption 8 (Uniform prior over next questions)**.** *Conditional on the feasible question set at stage k, the agent has no reason to favor any particular next question:*

$$\Pr(q = q_k^*) = \frac{1}{J_k}.$$

Assumption 9 (Shannon refinement costs)**.** *Refining beliefs over candidate next questions incurs Shannon (mutual-information) costs. If posterior beliefs are $q \in \Delta_{J_k}$ and the prior is π, the cost of refinement is*

$$\kappa D(q\|\pi) = \kappa \sum_{i=1}^{J_k} q_i \log\left(\frac{q_i}{\pi_i}\right).$$

Bellman problem

Let V_k denote the maximal expected continuation value when the agent is on the effective path at stage k.

Let

$$x(q) = \max_i q_i$$

denote the probability that the question selected after learning is the correct next question.

If the correct next question is asked the agent remains on the effective path and obtains continuation value V_{k+1}. If an incorrect question is asked the agent moves to an off-path state with continuation value V_{k+1}^{off}.

The Bellman problem is therefore

$$V_k = \max_{q \in \Delta_{J_k}} \Big[x(q) V_{k+1} + (1 - x(q)) V_{k+1}^{\text{off}} - \kappa D(q \| \pi) \Big].$$

A restricted learning problem

To solve the problem we again consider a restricted class of refinement policies that produce a single candidate next question.

Under such a policy the posterior takes the form

$$q(x) = \left(x, \frac{1-x}{J_k - 1}, \ldots, \frac{1-x}{J_k - 1} \right), \qquad x \in [1/J_k, 1].$$

Economically, this corresponds to allocating attention toward distinguishing one promising continuation from the remaining alternatives.

Restricted optimization

Under this posterior the Shannon cost becomes

$$C_{J_k}(x) = \kappa \left[x \log(J_k x) + (1-x) \log\left(\frac{J_k(1-x)}{J_k - 1} \right) \right].$$

The restricted Bellman problem becomes

$$V_k = \max_{x \in [1/J_k, 1]} \Big\{ x V_{k+1} + (1-x) V_{k+1}^{\text{off}} - C_{J_k}(x) \Big\}.$$

Proposition 3 (Optimal path discipline)**.** *The optimal probability of selecting the correct next question satisfies*

$$V_{k+1} - V_{k+1}^{\text{off}} = \kappa \log\left(\frac{x(J_k - 1)}{1 - x} \right).$$

Equivalently,

$$x_k^* = \frac{1}{1 + (J_k - 1)\exp\left(-\frac{V_{k+1} - V_{k+1}^{\text{off}}}{\kappa}\right)}.$$

Optimality of the restricted policy

As in the identification geometry, the payoff from any refinement policy depends only on the probability x of selecting the correct next question.

Lemma 2 (Cost-minimizing posterior)**.** *For any $x \in [1/J_k, 1]$, the posterior that assigns probability x to the chosen question and spreads the remaining probability uniformly over the other $J_k - 1$ questions uniquely minimizes the Shannon information cost.*

The proof follows directly from the strict convexity of relative entropy.

Theorem 4 (Optimal refinement under sequential discipline)**.** *Under Assumptions 7–9, the restricted refinement policy characterized above is globally optimal in the Bellman problem.*

Economic interpretation

Sequential-discipline geometry isolates a different planning margin from identification geometry. In identification environments the central challenge is recognizing which question resolves the uncertainty. In sequential-discipline environments the challenge is maintaining the correct trajectory of inquiry.

Optimal attention allocation therefore concentrates on preserving the effective path rather than identifying a decisive question. The optimal degree of discipline depends on the value difference between remaining on the effective path and stepping off it,

$$V_{k+1} - V_{k+1}^{\text{off}}.$$

When this value difference is large, attention is concentrated more sharply on identifying the correct next question. When the difference is small, deviations from the path are less costly and inquiry becomes less disciplined.

6.3 Comparison

Both geometries arise from the same Bellman recursion in Section 4. They differ only in diagnostic structure.

- In identification geometry, symmetry across remaining candidates coexists with strictly positive continuation value. Inquiry therefore persists until identification occurs.
- In sequential geometry, symmetry persists but continuation value collapses once the correct sequence is broken.

Diagnostic geometry alone—holding cost structure and horizon fixed— can therefore generate radically different inquiry dynamics.

6.4 From Invisible Planning to Observable Inquiry

The theory developed above isolates a distinct economic object: the organization of inquiry under uncertainty about diagnostic structure.

Historically, this object has been difficult to study empirically. The allocation of forward structural reasoning occurs largely before action, in internal deliberation, consultation, or informal exchange. The sequence of questions that precedes commitment typically leaves no observable trace. Differences in planning discipline are easily confounded with differences in priors, preferences, or luck.

For this reason, the model might appear, at first glance, to describe an invisible margin: rich in structure but difficult to verify.

That constraint is changing.

In environments where inquiry becomes explicit and recorded—most notably in interaction with artificial intelligence systems— questions are logged, answer sequences are preserved, and the structure of inquiry becomes observable.

When agents interact with AI systems, the questions are literal, the sequencing is traceable, and the forward exploration of contingencies can be reconstructed. The invisible planning margin becomes measurable.

The architecture developed in this section therefore acquires new empirical relevance. It provides a vocabulary for describing, measuring, and evaluating how inquiry is organized when questions are explicit.

The next section turns to this setting, examining how human–AI interaction makes the sequencing of questions observable, and how the theory of Question-Driven Learning can be brought to data.

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.

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.

One influential approach was protocol analysis. Participants in laboratory experiments were asked to "think aloud" while solving problems so that researchers could record the sequence of thoughts that occurred during the decision process.

Protocol analysis produced valuable insights. It revealed how individuals search for solutions, interpret evidence, and revise their beliefs. The method 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 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 Naturalistic Planning Logs

While protocol analysis attempted to observe reasoning in laboratory settings, recent technological developments make it possible to observe inquiry directly in natural 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. Each response may lead to new questions, revisions, or shifts in direction.

The resulting interaction logs capture the structure of inquiry in a way that was previously impossible.

Several features make these logs particularly valuable.

First, they are naturalistic. Individuals use AI systems to explore real decisions rather than artificial laboratory tasks.

Second, they are sequential. Each question appears in the context created by previous answers, allowing researchers to observe how inquiry evolves over time.

Third, they are incentivized. Users ask questions because they want something: a job, a strategy, clarification of uncertainty. The sequence of queries is therefore disciplined by real objectives rather than experimental instructions.

Taken together, these characteristics create an unprecedented opportunity to observe planning as it actually unfolds.

From model objects to observable traces. 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 about the payoff state and diagnostic structure before selecting the next question. Formally, the theory characterizes inquiry as a stochastic process over question–answer sequences

$$(q_1, a_1), (q_2, a_2), \ldots, (q_T, a_T),$$

generated by diagnostic refinement and state-dependent question choice.

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. As a result, the question–answer paths that constitute planning in the model become observable traces in real environments. The measurement problem therefore becomes one of interpreting these traces through the structure of the model.

7.4 From Transcripts to Measurement

Interaction logs record inquiry, 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.

A realized interaction trace

$$\mathcal{T} = \{(q_1, a_1), \ldots, (q_T, a_T)\}$$

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

The trace records ordering. It records branching. It records when inquiry stabilizes and when commitment occurs.

What it does not directly reveal are the internal beliefs or refinement processes that generated the trace. Interpreting inquiry logs therefore requires a model that links observable question sequences to latent diagnostic beliefs and refinement decisions.

Without such a model, logs are transcripts. With such a model, they become observable realizations of structured inquiry.

7.5 Natural Logs as Measurement Infrastructure

The emergence of large-scale AI interaction logs represents a structural transformation in the availability of economic data.

Millions of interactions occur daily. Individuals repeatedly engage AI systems for job search, skill acquisition, strategic planning, and adaptation to technological change. These traces are high-frequency, sequential, and persistent.

They reveal how individuals structure inquiry under uncertainty.

Researchers can observe the depth of exploration, the order in which questions appear, the moment when inquiry converges, and the point at which individuals stop asking further questions.

These patterns correspond closely to the components of diagnostic planning emphasized in the model.

Analyzing such data is not trivial. Diagnostic regimes vary across contexts, priors differ across individuals, and traces must be interpreted carefully. Nevertheless, the availability of sequential inquiry records creates a new empirical domain for the study of planning.

7.6 Model-Exact Experimental Designs

Natural logs provide rich observational data but embed uncontrolled variation in diagnostic structure and priors.

The model suggests a complementary approach: implementing environments that reproduce the structure of the dynamic inquiry problem directly.

In such settings the diagnostic regime D can be drawn from a known distribution and the experiment menu $\{\Pi_q^D\}$ explicitly defined. Participants may ask questions from a finite set $\mathcal{Q}$ and receive answers according to the underlying diagnostic structure.

Inquiry budgets can be limited to T steps and payoffs can depend on the probability that the participant identifies the correct decision.

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.

Under posterior-separable cost specifications, refinement intensity becomes a structural parameter that can be estimated from observed traces.

7.7 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.

7.8 Co-Evolution of Model and Measurement

Historically the absence of durable inquiry traces 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 traces can be interpreted. Conversely, the existence of such traces 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 traces 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 changes the economic environment precisely along this margin.

Economic models of information were developed for environments in which reliable information about the world was costly to obtain. Artificial intelligence alters this constraint. Modern AI systems generate candidate answers to an extraordinary range of questions almost instantly. Information that once required hours of search or consultation can now be produced within seconds.

When answers become inexpensive to generate, the scarcity that constrains decision-making shifts. The bottleneck moves from the acquisition of information to the cognitive effort required to organize inquiry.

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

8.1 Revisiting the Inquiry Problem

Consider again the worker exploring whether a transition into data science is feasible.

Using an artificial intelligence system, the worker can generate answers to a wide range of 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 about the environment. 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 data science field?
Is a specific credential typically required?
Given my background, how difficult would it be to obtain that credential?
What is the fastest credible way to acquire it?

In this sequence the investigation converges rapidly on the decisive constraint.

In the language of the model, the two inquiries correspond to different ways of navigating diagnostic structure. The first explores descriptive dimensions of the environment, while the second isolates the distinction that determines the decision.

Artificial intelligence therefore 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. Planning skill reflects several distinct capabilities in organizing inquiry.

First, individuals differ in the questions they are able to generate. The set of candidate questions considered by an agent is itself a cognitive object. Some planners quickly generate questions that reveal decisive distinctions, while others explore descriptive aspects of the environment without identifying the structural pivot of the problem.

Second, individuals differ in how they interpret answers. An answer may clarify the structure of the problem or suggest a new direction of investigation. Extracting these implications requires judgment.

Third, individuals differ in their ability to anticipate the diagnostic value of future questions. Skilled planners recognize which distinctions are likely to reorganize the investigation and direct inquiry accordingly.

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 on the diagnostic structure of the environment.

In relatively shallow environments, a single well-chosen question resolves most uncertainty. When answers can be generated quickly, artificial intelligence may substitute for routine cognitive effort.

In diagnostically complex environments, questions interact and early answers reshape the structure of inquiry. In such environments the ability to identify decisive distinctions and organize sequences of questions becomes more valuable.

Artificial intelligence therefore tends to complement planning skill in environments where inquiry structures are deep and sequential.

The framework developed in this paper provides a way to analyze this interaction formally. Differences in planning discipline correspond to differences in question sets, diagnostic beliefs, and the allocation of cognitive effort to evaluating which questions are informative.

8.4 Relation to Task-Based Frameworks

The dominant empirical framework for analyzing technological change in labor markets focuses on tasks.

Task-based models study how technologies substitute for or complement specific activities performed by workers. These frameworks have been extremely useful for understanding automation and changes in occupational structure.

The framework developed here highlights a different margin.

Rather than focusing on the execution of tasks, the model focuses on how agents organize inquiry before committing to action. Artificial intelligence systems do not merely automate specific tasks; they generate answers to questions. As a result, the productivity of decision-making increasingly depends on how inquiry is structured.

Planning under diagnostic uncertainty therefore operates at a different level of analysis than the task framework. The two perspectives are complementary rather than competing.

8.5 Measurement and Development

Historically the organization of inquiry was extremely difficult to observe.

Artificial intelligence systems change this constraint. Interaction logs preserve the sequence of questions through which decisions are explored. These traces reveal how individuals structure

inquiry under uncertainty.

Once planning becomes visible, it also becomes measurable. Planning capital therefore becomes observable.

Because the diagnostic structure of an environment can be specified in advance, the informational value of different questions can be evaluated directly. Participants who identify decisive distinctions increase the probability of reaching the correct decision, while those who pursue irrelevant lines of inquiry reduce it.

Performance can therefore be evaluated from the sequence of questions chosen by the participant rather than solely from the final outcome.

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 practice identifying decisive distinctions and organizing inquiry efficiently. Artificial intelligence systems provide a natural platform for such environments because they record inquiry sequences and allow rapid exploration of alternative strategies.

Planning capital therefore becomes not only measurable but teachable.

Artificial intelligence does not eliminate the need for planning. It makes the disciplined organization of inquiry both observable and developable.

From observable inquiry to economic analysis Artificial intelligence therefore alters more than the availability of information. It changes what can be studied.

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 be examined directly.

The framework developed in this paper provides a structure for interpreting those observations. It identifies the economic object underlying these inquiry traces: planning under diagnostic uncertainty. The conclusion returns to the broader implication of this shift.

9 Conclusion

Economic decisions are rarely made in a single step. They are preceded by inquiry.

Workers exploring career transitions investigate industries, training requirements, and geographic opportunities. Firms evaluating strategic investments examine markets, technologies, and competitors before committing resources. Researchers designing empirical tests explore alternative explanations and possible experiments.

In each of these environments, the quality of the final decision depends not only on the information available to the decision-maker but on how the investigation preceding action is organized.

Historically this margin remained largely invisible. Economists typically observed the outcomes of decisions rather than the reasoning that produced them. The sequence of questions through

which decisions were prepared rarely left a durable trace.

Artificial intelligence changes this constraint.

When answers become inexpensive to generate, the scarce input in decision-making becomes the cognitive effort required to organize inquiry effectively. The ability to generate decisive questions, interpret answers in ways that redirect investigation, and recognize when inquiry has converged becomes increasingly important.

The framework developed in this paper formalizes this margin. Planning under diagnostic uncertainty becomes a defined economic object rather than an informal intuition. Questions are treated as atomic acts, diagnostic regimes determine how answers should be interpreted, and agents allocate cognitive effort to evaluating which lines of inquiry are worth pursuing.

Artificial intelligence makes this margin visible. Interaction logs record sequences of questions and answers through which individuals explore decisions. Structured environments allow the informational value of different inquiries to be evaluated directly. Planning discipline can therefore be measured rather than inferred.

Once observable, it can also be developed. Environments that reward the identification of decisive distinctions allow individuals to practice structuring inquiry effectively.

The organization of inquiry is therefore no longer a residual category in economic analysis. It is a central component of decision-making in environments where answers are abundant but understanding remains scarce.

10 Intellectual Foundations and Related Traditions

The preceding sections developed a model of planning under diagnostic uncertainty. The framework introduced here—Question-Driven Learning—treats inquiry as a sequential process in which agents select questions before committing to consequential actions. Planning consists of organizing that inquiry when the diagnostic value of available questions is itself uncertain.

This section situates the framework within several existing intellectual traditions. These literatures illuminate different aspects of the problem studied in this paper: the construction of problem representations, heterogeneity in planning ability, formal models of information acquisition, and the measurement of reasoning processes.

The goal is not to replace these traditions but to clarify how the present framework extends them. Question-Driven Learning identifies a distinct economic margin: the endogenous sequencing of constrained questions under uncertainty about what questions reveal prior to irreversible commitment.

10.1 Problem Construction and the Legacy of Herbert Simon

The intellectual roots of this framework lie in Herbert Simon's theory of bounded rationality (Simon, 1947, 1996; Newell and Simon, 1972). Simon emphasized that decision-making rarely begins with a

fully articulated set of alternatives. Instead, individuals construct problem representations through sequential refinement.

In Simon's analysis of means–ends reasoning, alternatives are generated and decomposed through structured inquiry (Newell and Simon, 1972). The agent's task is therefore not merely to select among known options but to organize a complex environment into tractable distinctions under cognitive constraint.

Two implications follow.

First, the representation of the problem is endogenous. The distinctions considered, the contingencies recognized, and the criteria used to evaluate alternatives emerge through a process of sequential reasoning.

Second, performance depends on the structure of that reasoning. Different sequences of decomposition may expose or conceal decisive distinctions. Economic divergence may therefore arise before any explicit choice is made.

Simon's framework placed this interrogative layer at the center of reasoning. However, it did not embed interrogative sequencing as a formally specified economic choice variable within a dynamic payoff environment. Planning was described heuristically rather than modeled as a constrained economic decision.

Question-Driven Learning translates Simon's insight into such an environment. The framework treats interrogative sequencing—deciding which question to clarify next under diagnostic uncertainty—as a payoff-relevant economic choice.

10.2 Cognitive Economics and Structured Mistakes

If Simon established that problem representation is endogenous, Cognitive Economics advances a complementary claim: cognitive processes themselves are economically consequential objects that require formal specification and measurement (Caplin, 2023, 2016).

Rather than treating deviations from full rationality as residual noise, Cognitive Economics studies the structured mechanisms through which attention is allocated, distinctions are selected, beliefs are revised, and mistakes propagate into economic outcomes.

Two methodological commitments define this program (Caplin, 2024). First, cognitive constraints must be embedded in explicit payoff environments. Second, theoretical constructs must be tied to observable data.

Within this framework, planning failures are not accidental. Agents may misidentify decisive distinctions, allocate inquiry toward low-diagnostic contingencies, or terminate refinement prematurely. Such patterns can generate persistent divergence even when individuals share similar preferences and resources.

Question-Driven Learning operationalizes this perspective. The model specifies how agents allocate attention across potential questions, how answers reshape beliefs about diagnostic structure, and how inquiry sequences influence final decisions. Planning skill therefore becomes a defined economic margin rather than an informal intuition.

10.3 Planning in Cognitive and Developmental Psychology

Cognitive and developmental psychology has long studied planning ability directly. Experimental paradigms isolate sequential reasoning, multi-step foresight, and executive control in structured problem-solving environments (Newell and Simon, 1972; Ericsson and Simon, 1980).

These studies reveal substantial heterogeneity in planning ability. Individuals differ in their capacity to sequence intermediate steps, anticipate downstream consequences, and revise strategies as new information becomes available.

Two features of this literature are particularly relevant.

First, planning ability is not reducible to general intelligence. Individuals with similar cognitive test scores may differ substantially in their ability to structure multi-step reasoning.

Second, planning is inherently sequential. Performance depends on the order in which subproblems are addressed and on how inquiry evolves over time.

However, the environments used to measure planning have traditionally been laboratory tasks. While these experiments demonstrate heterogeneity in planning ability, they rarely embed planning within economically consequential decisions.

Interactive AI environments alter this constraint. Planning now unfolds through explicit question–answer exchanges in domains such as job search, retraining, and strategic decision-making. Sequential traces therefore preserve the structure of inquiry in environments where decisions carry real economic consequences.

10.4 Planning Skill and the Economics of Human Capital

The economics of skill has progressively expanded its scope. Early work emphasized general cognitive ability—often proxied by IQ—as a determinant of educational attainment and earnings. Subsequent research incorporated broader conceptions of human capital, including persistence, self-regulation, and social interaction (Heckman et al., 2006; Deming, 2022).

Task-based models later reframed skill in terms of execution (Autor et al., 2003; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2020). Occupations were decomposed into activities, and technological change was analyzed through its impact on task productivity.

These frameworks illuminate important dimensions of economic behavior. However, they focus primarily on execution within an articulated problem structure.

Planning skill concerns an earlier stage. It governs how that structure is constructed before commitment.

Agents must determine which distinctions are relevant, which contingencies merit investigation, and when inquiry has converged sufficiently to support a decision. Two individuals with similar education and similar access to technology may therefore diverge systematically if they differ in how effectively they organize inquiry before acting.

10.5 Formal Models of Information Acquisition

Economic theory provides several rigorous frameworks for analyzing learning before commitment.

Search theory models how agents acquire information from known distributions (Mortensen and Pissarides, 1994). Bandit models study exploration across known experiments (Gittins, 1979; Berry and Fristedt, 1985). Rational inattention models analyze how agents allocate attention across signals about payoff states (Sims, 2003; Matejka and McKay, 2015; Caplin and Dean, 2015).

Across these literatures the informational structure of the environment is typically specified ex ante. Agents optimize over intensity—how long to search, how much precision to allocate, or which experiment to sample next.

Question-Driven Learning introduces a different margin. Here the informational value of questions themselves is 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.

The object of choice becomes the sequencing of questions that progressively clarifies the structure of the decision problem.

10.6 Measurement and Data Engineering

Historically the organization of inquiry was extremely difficult to observe. Economic data typically record outcomes rather than the reasoning that produced them.

Interactive AI systems change this constraint. When individuals explore decisions through question–answer exchanges with AI systems, the structure of inquiry becomes observable.

Sequential interaction traces record the questions asked, the answers received, and the moment when inquiry converges or stops.

These traces provide the empirical foundation for measuring planning behavior. The framework developed in this paper interprets such traces as realizations of a dynamic inquiry process under diagnostic uncertainty.

Measurement therefore becomes possible precisely because inquiry leaves a durable record.

10.7 Artificial Intelligence and the Upstream Margin of Adaptation

Much of the economic debate surrounding artificial intelligence has focused on tasks. Task-based models analyze how automation alters the execution of specific activities (Autor et al., 2003; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2020).

The framework developed here highlights a different margin.

Artificial intelligence drastically reduces the cost of generating answers while leaving the direction of inquiry to the user. As answer costs fall, the disciplined organization of inquiry becomes increasingly important.

Differences in planning skill may therefore influence economic outcomes even when individuals have similar access to technology. Divergence may arise before any task is executed or investment made, in the organization of inquiry that precedes commitment.

Question-Driven Learning formalizes this upstream layer. It identifies planning under diagnostic uncertainty as a distinct economic margin that becomes particularly consequential in environments where answers are abundant but attention 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.

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. and Fristedt, B. (1985). *Bandit Problems: Sequential Allocation of Experiments*. Chapman and Hall.

Caplin, A. (2016). *The Science of Mistakes*. World Scientific.

Caplin, A. (2023). *An Introduction to Cognitive Economics: The Science of Mistakes*. Palgrave Macmillan.

Caplin, A. (2024). Data engineering for cognitive economics. *Journal of Economic Literature*. Update volume and page numbers.

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.

Deming, D. (2022). Four facts about human capital. *Journal of Economic Perspectives*, 36(3):75–102.

Ericsson, K. A. and Simon, H. (1980). Verbal reports as data. *Psychological Review*, 87(3):215–251.

Gittins, J. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society: Series B*, 41(2):148–177.

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.

Mortensen, D. and Pissarides, C. (1994). Job creation and job destruction in the theory of unemployment. *Review of Economic Studies*, 61(3):397–415.

Newell, A. and Simon, H. (1972). *Human Problem Solving*. Prentice Hall.

Savage, L. J. (1954). *The Foundations of Statistics*. John Wiley & Sons, New York.

Simon, H. (1947). *Administrative Behavior*. Macmillan.

Simon, H. (1996). *The Sciences of the Artificial*. MIT Press.

Sims, C. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.