When Learning Increases Uncertainty: A Model of Multidimensional Consumer Evaluation

W. Jason Choi[*] Jiwoong Shin[†]

February 2026

Please do not circulate

Abstract

Consumers often learn about products in environments where uncertainty is multidimensional: they face two orthogonal sources of uncertainty, about attribute levels (what the product delivers) and about attribute importance (how attributes map into value). In such settings, learning can have a counterintuitive consequence. Resolving one uncertainty can increase overall uncertainty in the consumer's valuation by amplifying the relevance of the remaining unknown dimension. We develop a Bayesian model of consumer evaluation in which utility is a weighted combination of two attributes, where both the weight and at least one attribute level are initially uncertain. The consumer can acquire information about either the weight or the uncertain attribute at a cost before deciding whether to purchase. In benchmark environments with only attribute uncertainty, information acquisition monotonically reduces utility uncertainty. In contrast, with both weight and attribute uncertainty, equilibrium learning can increase ex post uncertainty for intermediate information costs, even though each learning action is individually informative. We interpret this mechanism as a rational foundation for a Dunning–Kruger type pattern in consumer confidence: a coarse early assessment can feel decisive, while further learning can reveal that what remains unknown is precisely what matters most, lowering confidence. We discuss implications for information design, including what to disclose first and when transparency can backfire.

Keywords: Multidimensional uncertainty, information acquisition, Bayesian learning, variance decomposition, rational inattention, consumer search

[*]Robert H. Smith School of Business, University of Maryland; wjchoi@umd.edu

[†]School of Management, Yale University; jiwoong.shin@yale.edu.

1 Introduction

Consumers often evaluate multi-attribute products that are difficult to assess. Consider a first-time buyer choosing an electric vehicle (EV). She can readily compare salient attributes such as price or advertised range, but other dimensions (e.g., charging speed in cold weather, battery degradation, repair convenience, or software reliability) are harder to interpret ex ante. Moreover, she may be uncertain about which of these attributes will ultimately matter most during usage. A test drive or additional research may resolve some experiential features, yet also redirect attention toward hard-to-forecast factors. Paradoxically, the more she learns, the less confident she may become about the product's overall utility.

At first glance, this pattern is puzzling. In standard learning models where uncertainty lies along a single payoff-relevant dimension, additional information weakly reduces posterior dispersion: more information should clarify the decision and increase confidence. If the buyer were uncertain only about a single latent "overall quality," further reviews or testing would monotonically reduce uncertainty. In practice, however, consumers often face multidimensional uncertainty, not only about attribute levels, but also about how those attributes map into value. As learning proceeds, consumers may discover that their evaluation hinges on dimensions whose importance is itself uncertain. In such environments, information can shift uncertainty toward what matters most, potentially lowering confidence, expanding search, and delaying purchase.

This paper provides a rational, micro-founded account of this non-monotonic confidence pattern. Without invoking behavioral biases or mistaken updating, we show that product-utility uncertainty can be non-monotonic in the consumer's learning cost; this pattern emerges when uncertainty spans multiple, interacting dimensions. The key insight is that learning in multidimensional environments can *reallocate* uncertainty across orthogonal dimensions, rather than simply eliminate it. This amplification mechanism does not arise when uncertainty is confined to a single dimension.

We formalize this idea in a Bayesian model of product evaluation. A consumer's valuation depends on two attributes combined through a relative-importance weight. One attribute is salient and easy to assess; the other is harder to evaluate ex ante. The consumer is uncertain both about the level of the hard-to-assess attribute and about the weight that attribute will receive in her realized utility. Before choosing between the product and an outside option, she can pay a cost to learn either (i) which attribute will ultimately be decisive, or (ii) the realization of the uncertain attribute itself. This framework allows us to ask: can fully rational Bayesian learning make a consumer less certain about her valuation?

The answer depends on the dimensionality of uncertainty. When only attribute levels are uncertain and weights are known, learning has the standard effect of reducing valuation uncertainty. When uncertainty concerns both attribute levels and attribute importance, however, learning one dimension can increase the exposure of valuation to the other, unresolved dimension. Resolving weight uncertainty, for example, may reveal that valuation hinges on a still-uncertain attribute level; conversely, learning an attribute realization can make the uncertain weight more consequential. Even with correct Bayesian updating, additional information can therefore increase ex post uncertainty in some states: not because beliefs are distorted, but because learning shifts uncertainty toward what is payoff-relevant.

We then characterize the consumer's optimal information acquisition strategy when learning is costly. If learning is cheap, the consumer learns both dimensions and uncertainty collapses; if it is costly, she learns neither and uncertainty remains unchanged. For intermediate learning costs, however, the consumer optimally resolves only one dimension. In this case, learning can increase ex post valuation variance in states where the remaining uncertainty becomes especially consequential. Thus, deeper search may reduce confidence even as it adds information.

We interpret this mechanism as a rational analogue of the Dunning–Kruger pattern (Kruger and Dunning, 1999). Consumers who stop early in the learning process may appear confident because they do not uncover how consequential the unresolved dimension is. Consumers who learn more can become less confident because they discover that what matters most is precisely what remains uncertain. In extensions, we show how heterogeneity in information costs generates systematic differences in learning intensity and reported confidence: high-cost consumers acquire coarse information and appear more confident, whereas low-cost consumers acquire richer information and may report lower confidence after learning.

Beyond its behavioral interpretation, the mechanism has implications for information design.

Firms and platforms choose what to disclose and in what sequence. Disclosure about attribute levels and disclosure about attribute importance are not interchangeable: sequencing shapes how remaining uncertainty affects valuation. Revealing weight information first can shift attention toward uncertain attributes, while revealing attribute realizations first can magnify the consequences of weight uncertainty. These insights offer a new perspective on staged disclosure, transparency tools, and comparison platforms.

The remainder of the paper develops these results in detail. We discuss related literature in Section 2, and introduce the basic model and information structure in Section 3. Section 4 builds intuition by comparing one-dimensional and two-dimensional uncertainty and derives the key variance amplification channel. In Section 5, we characterize equilibrium learning and identify parameter regions in which learning increases ex post uncertainty. Section 6 provides extensions, including continuous signal precision choice and endogenous attention, and discusses implications for firm information provision and disclosure sequencing. Section 7 concludes.

2 Literature Review

Our paper relates to several streams of research in marketing and economics on consumer learning, information acquisition, and multiattribute evaluation. A large literature models consumers as Bayesian learners who update beliefs about product quality over time. Early work such as Erdem and Keane (1996a) and Erdem and Keane (1996b) embeds learning about latent quality within dynamic discrete choice frameworks, showing how posterior beliefs evolve with experience and how learning shapes brand choice. Subsequent research extends this approach to richer state-space representations and dynamic brand equity models (e.g., Sriram et al. 2006; Sriram and Kalwani 2007). In these environments, uncertainty typically lies in a single latent quality dimension, and learning monotonically reduces posterior variance. Our framework differs by allowing uncertainty to span two orthogonal dimensions: uncertainty over attribute realizations and uncertainty over the weights that map attributes into utility. When these dimensions interact, learning can reallocate uncertainty across them rather than uniformly compress it.

Our analysis relates to a broad literature on costly information acquisition and deliberation in marketing. Beginning with Shugan (1980), researchers model information processing as costly and show that consumers optimally allocate attention across product attributes. Subsequent work develops structural models in which consumers endogenously choose which information to acquire and how intensively to search. For example, Guo and Zhang (2012) analyze how consumers optimally acquire attribute information and how firms adjust pricing in response; Branco et al. (2016) examine equilibrium information provision when search costs are endogenous; and Kuksov and Villas-Boas (2010) show that expanding alternatives or attributes can reduce decision quality even under fully rational behavior. Related work studies how firms design disclosure policies, product lines, and pricing schemes when consumers face costly deliberation (e.g., Li et al. 2019; Li and Hitt 2021; Liu and Dukes 2022; Dukes and Liu 2016). Across these models, information acquisition is endogenous and costly, and consumers may rationally ignore some dimensions when processing costs are high.

A formal theoretical foundation for this approach is provided by rational inattention models (Sims, 2003; Matejka and McKay, 2015; Caplin and Dean, 2015), which characterize optimal signal acquisition subject to information constraints. In standard one-dimensional environments, additional information is Blackwell more informative and therefore weakly reduces posterior dispersion. Our framework departs from this benchmark by allowing uncertainty to span interacting dimensions of the valuation problem. When consumers are uncertain both about attribute realizations and about the weights that map attributes into utility, resolving one dimension can increase the payoff sensitivity to the other. Consequently, optimal learning can increase ex post valuation uncertainty in equilibrium, even though each learning action is individually informative.

Our framework is closely related to the multiattribute evaluation literature. Classical models treat utility as a weighted combination of attributes and study how consumers form consideration sets and trade off product characteristics (e.g., Hauser and Wernerfelt 1990). Also, Dzyabura and Hauser (2019) examine environments in which consumers learn their own preference weights and firms design recommendation systems accordingly. Their analysis highlights the economic importance of weight uncertainty, but learning reduces uncertainty in the conventional sense and does not generate interaction effects across uncertainty dimensions. A particularly relevant recent contribution is Lee et al. (2025), who study how firms strategically communicate attribute importance under competition when consumers are uncertain about which of two main attributes matters more. In their setting, weight uncertainty interacts with strategic disclosure and competitive incentives, but attribute realizations are typically taken as given. We extend this line of work by allowing simultaneous uncertainty over both weights and attribute realizations, and by showing that interaction between these uncertainties fundamentally alters the effect of learning on valuation dispersion.

Finally, our work connects to research on consumer confidence and evaluative dynamics under uncertainty. A substantial literature documents that confidence and belief accuracy do not always move in parallel (e.g., Moore and Healy 2008; Puntoni et al. 2009). In marketing contexts, studies show that additional information can alter confidence in systematic ways, sometimes increasing doubt rather than resolving it (e.g., Lee and Ariely 2011, Tormala and Petty 2006). Related research on search and decision processes suggests that deeper information acquisition can heighten awareness of tradeoffs and increase perceived ambiguity (e.g., Carmon et al. 2003, Diehl and Poynor 2004). Much of this work emphasizes cognitive limits, biased self-assessment, or psychological mechanisms behind overconfidence and confidence revision. In contrast, we provide a fully rational account in which confidence can decline even under correct Bayesian updating. In this sense, our model offers a rational, multidimensional foundation for confidence patterns often associated with the Dunning–Kruger effect (Kruger and Dunning, 1999).

3 Basic Model

The game consists of a consumer and a firm. The consumer has unit demand for the firm's product and chooses whether to purchase it or take an outside option normalized to zero. Throughout the paper, we use the terms *utility* and *valuation* interchangeably to denote the consumer's perceived payoff from the product. Utility should be interpreted broadly as overall product value, which may reflect quality, performance, risk, or other payoff relevant dimensions. We abstract from pricing in the baseline model to isolate informational forces; introducing an additive price term would not affect the mechanism developed below.

The product is characterized by two payoff-relevant attributes. The first attribute, denoted α, is relatively salient and easier to assess prior to purchase. The second attribute, denoted β, is harder to pin down ex ante and is learned only imperfectly from pre-purchase information. For example, consider a first-time EV buyer. She can usually form a fairly crisp assessment of a salient dimension such as advertised driving range or a well-communicated performance metric, which we bundle into α. By contrast, other performance dimensions are intrinsically harder to forecast before living with the car, such as charging speed in cold weather, battery degradation, or the reliability of software and charging access in her day-to-day routine, which we bundle into β.

$$u \;=\; \omega\alpha + (1-\omega)\beta, \tag{1}$$

where $\omega \in [0,1]$ is an *importance weight* that governs the relative importance of these two attributes in determining utility. Thus, it is an unknown payoff relevance primitive that (i) affects how utility is generated and (ii) can be learned separately from the uncertain attribute level β.

A key modeling ingredient is that the consumer is uncertain not only about *what* the product delivers (the realization of β), but also about *how* that delivery will matter for her payoff (the value of ω). In the EV example, ω captures how much the buyer's realized satisfaction will end up hinging on the easy-to-assess dimension α versus the hard-to-forecast dimension β once she starts using the vehicle. This importance weight can be interpreted as summarizing the consumer's future usage environment and the tradeoffs she will face after adoption, such as whether she has reliable home charging, how often she takes long trips, or how costly it is for her to handle charging uncertainty.[1]

In the baseline specification, we assume that α is known, while β and ω may be uncertain from the consumer's perspective. This structure allows us to focus on two dimensions of uncertainty: uncertainty regarding the attribute β and uncertainty over how strongly β will matter for utility. The next subsection makes this distinction precise by separating uncertainty about outcomes from uncertainty about how outcomes map into payoff.

[1] The model does not require ω to be purely product-specific; it can be consumer-specific or context-specific. What matters for our mechanism is that, at the time of evaluation, ω is a payoff-relevant latent state that is not yet known to the consumer.

Two orthogonal sources of uncertainty

A central feature of the model is that the consumer faces two conceptually distinct, and a priori separable, sources of uncertainty. The first is outcome uncertainty: the consumer is unsure about what the product will deliver along at least one payoff relevant dimension. In the baseline model, this uncertainty is captured by the attribute β, which is difficult to assess before purchase. Continuing the EV example from the introduction, β can be read as a performance dimension that is learned only imperfectly from brief inspection or a short test drive, such as charging speed in cold weather, battery degradation over time, repair convenience, or software reliability. We summarize the consumer's prior uncertainty about this attribute by its mean and variance,

$$\mathbb{E}[\beta] = \mu_\beta, \qquad \mathrm{Var}(\beta) = \sigma_\beta^2,$$

and later specify how information acquisition changes the posterior distribution of β.

The second is mapping uncertainty: the consumer is unsure about how attribute realizations translate into her realized payoff. In practice, this reflects uncertainty about which dimension will end up being decisive once the consumer starts living with the product, which can depend on future usage conditions and on what constraints become binding. In the EV example, a consumer may initially focus on an easily comparable dimension such as advertised range or comfort, but after learning a bit she may realize that charging access or battery performance will be what ultimately drives her experience, even though those dimensions are harder to forecast. In the model, we capture this mapping uncertainty with a weight parameter $\omega \in [0, 1]$ in the utility specification in (1), where $u = \omega\alpha + (1 - \omega)\beta$.

Higher realizations of ω correspond to states in which the salient, easy to assess dimension α is relatively more important for realized utility; lower realizations correspond to states in which the hard to assess dimension β is relatively more decisive. We summarize the consumer's prior uncertainty about the attribute importance parameter by

$$\mathbb{E}[\omega] = \bar{\omega}, \qquad \mathrm{Var}(\omega) = \sigma_\omega^2,$$

The term *orthogonal* is meant in an informational sense: learning about the attribute realization β need not directly resolve uncertainty about the mapping parameter ω, and learning about ω need not directly resolve uncertainty about β. To isolate the amplification mechanism cleanly, we take ω and β to be independent under the prior unless otherwise stated. This rules out correlation driven effects and ensures that any non monotonicity in uncertainty comes from the interaction created by the utility mapping itself.

It is important to note that this separation between outcome uncertainty and mapping uncertainty is broader than the specific interpretation of "attribute level" versus "attribute importance." More generally, the first dimension can be understood as uncertainty about payoff relevant primitives, while the second dimension can be interpreted as uncertainty about the payoff function that maps those primitives into realized utility. The analysis below shows that when both dimensions are present, learning can shift where uncertainty matters rather than simply reduce it.

3.1 Information acquisition

Before choosing whether to purchase the product, the consumer can acquire information about the two primitives that are initially uncertain from her perspective: the uncertain attribute level β (outcome uncertainty) and the weight parameter ω (mapping uncertainty). We interpret information acquisition broadly to include deliberate search (reading reviews, consulting experts), product experience (a test drive), or platform mediated diagnostics that clarify either what the product delivers or which dimension will be important in the consumer's realized payoff.

The consumer can acquire at most one signal about each dimension. Let $d_\omega \in \{0, 1\}$ indicate whether the consumer acquires information about ω, and let $d_\beta \in \{0, 1\}$ indicate whether she acquires information about β. Acquiring information is costly: learning ω costs $c_\omega > 0$ and learning β costs $c_\beta > 0$. In the baseline model we set $c_\omega = c_\beta = c > 0$ to focus attention on the interaction between uncertainty dimensions; allowing asymmetric costs is straightforward and useful for comparative statics.

Signals

Tthe baseline analysis assumes that learning is perfectly revealing conditional on acquisition (which will be relaxed later):

$$s_\omega = \omega \quad \text{if } d_\omega = 1, \qquad s_\beta = \beta \quad \text{if } d_\beta = 1. \tag{2}$$

If a signal is not acquired, it is absent from the consumer's information set. Because the model is Bayesian, whenever a signal is observed the consumer updates beliefs using Bayes' rule. Under the independence assumption between ω and β, learning one dimension does not update the posterior over the other; however, learning one dimension can still change how consequential the remaining uncertainty is for her utility, which is the core force in the paper.

Purchase decision and information acquisition

After any information acquisition, the consumer decides whether to purchase the product or take an outside option normalized to zero, $u_0 = 0$. Let I denote the information set at the purchase stage. The consumer is risk neutral and chooses the action that maximizes her expected payoff:

$$a(I) \in \underset{a\in\{0,1\}}{\operatorname{argmax}} \ a \cdot \mathbb{E}[u|I], \tag{3}$$

so that the product is purchased whenever $\mathbb{E}[u|I] \geq 0$. It is convenient to define the value after observing information I as

$$V(I) \ = \ \max\{0, \mathbb{E}[u|I]\}. \tag{4}$$

Information affects decisions by changing posterior beliefs about u, and therefore the induced value $V(I)$.

Ex ante, before observing any signals, the consumer chooses a learning strategy to maximize expected value net of information costs. Let

$$\mathcal{D} = \{\emptyset,\ \{\omega\},\ \{\beta\},\ (\omega \to \beta),\ (\beta \to \omega)\}$$

denote the set of feasible learning strategies. These correspond to learning nothing ($\emptyset$), learning only one dimension (i.e., either $\{\omega\}$ or $\{\beta\}$), or learning both dimensions in either order (($\omega \to$

β), ($\beta \to \omega$)).

The consumer's problem is

$$\max_{d \in \mathcal{D}} \mathbb{E}\left[V(I(d))\right] \; - \; C(d), \tag{5}$$

where $I(d)$ is the information set induced by strategy d, and $C(d)$ denotes the total information cost under strategy d, so that $C(d)$ equals c_ω if only ω is learned, c_β if only β is learned, and $c_\omega + c_\beta$ if both are learned. Note that sequential learning is allowed under the strategies ($\omega \to \beta$) and ($\beta \to \omega$), where the second acquisition decision may depend on the realization of the first signal. The formulation in (5) therefore encompasses both one-shot and sequential learning without requiring explicit dynamic programming formulation.[2]

3.2 Baseline discrete model

To obtain sharp characterizations of equilibrium learning behavior, we first analyze a discrete baseline model. This specification allows us to isolate the interaction between outcome uncertainty and mapping uncertainty in the simplest possible environment. The main mechanism does not depend on discreteness; later sections relax these assumptions by allowing continuous primitives and noisy signals.

Outcome uncertainty

The attribute α is known and constant:

$$\alpha \in \mathbb{R}, \qquad \mathrm{Var}(\alpha) = 0. \tag{6}$$

This captures a salient, easily assessed dimension, such as a transparent specification or a familiar performance metric.

On the other hand, the attribute β is uncertain and takes two values:

$$\beta \in \{\beta_L, \beta_H\}, \qquad \mathbb{P}(\beta = \beta_L) = \mathbb{P}(\beta = \beta_H) = \frac{1}{2}, \tag{7}$$

[2]For example, if the consumer first chooses $d_\omega \in \{0, 1\}$, V denotes the consumer's ex ante value,

$$V = \max_{d_\omega \in \{0,1\}} \left\{ -c_\omega d_\omega + \mathbb{E}\left[\max_{d_\beta \in \{0,1\}} \{-c_\beta d_\beta + \max\{0, \mathbb{E}[u|I]\}\} \right] \right\},$$

where expectations are taken over realizations of (ω, β) consistent with the chosen strategy.

where $\beta_L = \mu_\beta - \sigma_\beta$ and $\beta_H = \mu_\beta + \sigma_\beta$, so that $\mathbb{E}[\beta] = \mu_\beta$ and $\mathrm{Var}(\beta) = \sigma_\beta^2$. This dimension represents a harder to forecast payoff component that is often learned only through deeper engagement.

Mapping uncertainty

The attribute weight parameter is also assumed to be binary:

$$\omega \in \{0, 1\}, \qquad \mathbb{P}(\omega = 1) = 1 - \mathbb{P}(\omega = 0) = q, \tag{8}$$

so that $\bar{\omega} = q$ and $\mathrm{Var}(\omega) = q(1 - q)$. When $\omega = 1$, utility is driven entirely by α; when $\omega = 0$, utility is driven entirely by β. This is a stark way to capture mapping uncertainty as uncertainty about which dimension will turn out to be decisive in the consumer's realized payoff. Moreover, unless otherwise stated, we assume ω and β to be independent to isolate the variance amplification channel from any mechanical correlation between outcome and mapping uncertainty.

Game sequence

The sequence of events is as follows. First, nature draws (ω, β) from the prior distribution. The consumer then chooses whether to acquire information about one of the two uncertainty dimensions at cost c, and observes the corresponding signal if she does so. After observing any acquired signal, she may acquire information about the remaining dimension, again at cost c. Finally, given her resulting information set, she decides whether to purchase the product or take the outside option.

We solve for subgame perfect equilibrium in the consumer's learning strategy. Because learning is sequential and the value of learning the second dimension depends on the realization of the first signal, the order of information acquisition is endogenous. This dependence generates equilibrium regions in which the consumer optimally resolves only one source of uncertainty, which is central to the mechanism developed below.

4 Uncertainty Decomposition and Interaction

We begin by characterizing how information affects uncertainty about utility, holding learning behavior fixed. This section isolates the core mechanism of this research by contrasting environments with one-dimensional and multidimensional uncertainty.

4.1 Benchmark: One-dimensional uncertainty

We begin by clarifying the object of interest. In addition to expected utility, we track the consumer's perceived uncertainty about the product through the posterior variance of utility, $\mathrm{Var}(u|I)$, conditional on the information set I. We interpret posterior dispersion as an operational measure of perceived confidence: a higher $\mathrm{Var}(u|I)$ indicates a wider range of plausible utility realizations consistent with the information observed, and therefore lower confidence in the product's payoff. Because the purchase decision depends only on the posterior mean $\mathbb{E}[u|I]$, this dispersion measure complements the mean-based decision rule by capturing how confident the consumer perceives the choice to be.

Consider first an environment in which uncertainty lies only in the outcome dimension. Suppose ω is known and only β is uncertain. Utility is then

$$u = \omega\alpha + (1-\omega)\beta,$$

and uncertainty arises solely from β. In this case,

$$\mathrm{Var}(u) = (1-\omega)^2\sigma_\beta^2.$$

Information about β reduces σ_β and therefore reduces $\mathrm{Var}(u)$. More informative signals in the sense of Blackwell[3] weakly reduce posterior variance. Thus, when uncertainty is one-dimensional, learning monotonically reduces perceived uncertainty about utility. This monotonicity property serves as the benchmark against which the multidimensional case can be compared.

4.2 Multi-dimensional uncertainty

We now consider the case in which uncertainty lies in both dimensions: the outcome β and the mapping parameter ω are both uncertain. Recall that $u = \omega\alpha + (1-\omega)\beta$. It is useful to rewrite

[3] A signal is more informative in the Blackwell (1953) sense if it induces a posterior that is a mean-preserving contraction relative to a less informative signal. In such environments, posterior dispersion weakly decreases. See DeGroot (1970) for a decision-theoretic treatment.

utility as

$$u = \beta + \omega(\alpha - \beta),$$

which makes explicit that utility depends on both the realization of β and the interaction between ω and the attribute gap $(\alpha - \beta)$. Under independence between ω and β, utility variance can be obtained directly

$$\mathrm{Var}(u) = (1-q)\sigma_\beta^2 + q(1-q)(\alpha - \mu_\beta)^2. \tag{9}$$

The first term reflects exposure to outcome uncertainty, while the second term reflects exposure to mapping uncertainty and depends on both σ_ω^2 and σ_β^2.

4.3 Non-monotonic learning effects

Equation (9) reveals the key interaction. In contrast to the one-dimensional benchmark, the contribution of mapping uncertainty depends on both σ_ω^2 and σ_β^2, as well as on the gap $(\alpha - \mu_\beta)^2$. Thus, the economic relevance of one uncertainty dimension depends on the state of the other.

To visualize this contrast, Figure 1 compares the one-dimensional benchmark with the multi-dimensional case.

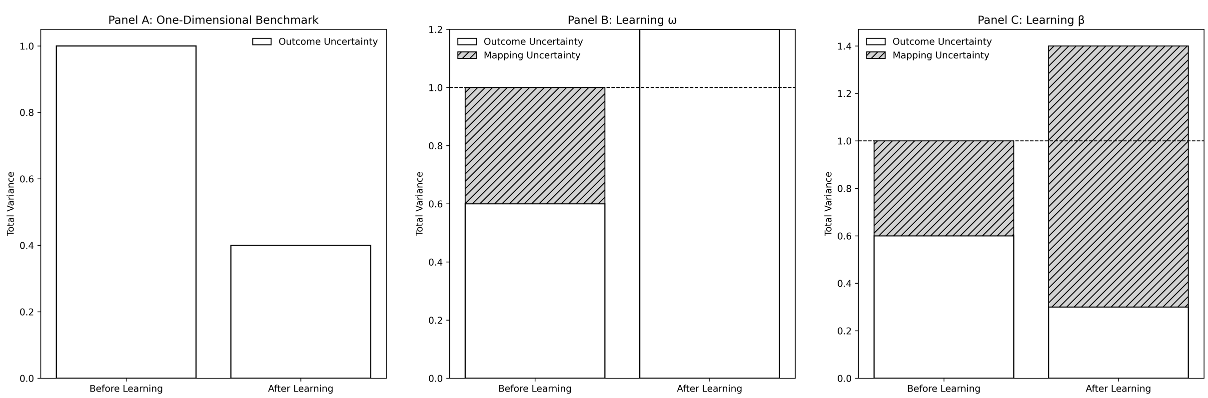

(a) One-dimensional benchmark (b) Learning ω (c) Learning β

Figure 1: Uncertainty under one- and two-dimensional environments. In each panel, the lighter component represents outcome uncertainty and the darker component represents mapping uncertainty. Panel A shows the benchmark case in which only β is uncertain and learning always reduces total variance. Panels B and C illustrate that, under multidimensional uncertainty, resolving one dimension can increase total utility variance by increasing exposure to the remaining uncertain dimension.

Panel A illustrates the benchmark case: when only β is uncertain, learning monotonically reduces total variance.[4] Panels B and C illustrate the multidimensional case.

Consider first learning the mapping parameter ω (Panel B). If the realized ω implies that utility is primarily driven by β, then outcome uncertainty becomes fully exposed in utility. Although learning the mapping parameter ω eliminates mapping uncertainty σ_ω^2, it increases the effective weight placed on outcome uncertainty σ_β^2. In some realizations, total variance after learning can exceed its pre-learning level. Conversely, consider learning the attribute realization β (Panel C). If the realized β differs substantially from α, the gap $(\alpha - \beta)^2$ increases, amplifying the contribution of mapping uncertainty to overall variance. Thus, resolving outcome uncertainty can increase the economic relevance of σ_ω^2. In both cases, resolving one dimension can increase total utility dispersion in some realizations.

Importantly, learning does not mechanically increase uncertainty. Figure 1 illustrates that amplification can arise for certain realizations and parameter configurations, but not universally. The central insight is that when uncertainty spans orthogonal dimensions, learning can shift where uncertainty matters rather than simply compress it. This interaction is absent in one-dimensional environments, where information monotonically reduces dispersion. By contrast, in the multidimensional setting the interaction between outcome and mapping uncertainty can make the remaining dimension economically more consequential after learning. In the next section, we show that this mechanism can generate equilibrium regions in which ex post utility variance exceeds ex ante variance.

5 Equilibrium Learning in Discrete Baseline

We now characterize the consumer's equilibrium information acquisition behavior in the discrete baseline model and ask a central question of the paper: can fully rational learning increase ex post uncertainty about utility along an equilibrium path? Section 4 established that resolving one uncertainty dimension can increase utility variance for some realizations. Here we show that this force survives endogenous information acquisition. We proceed in three steps. We first characterize

[4] This benchmark relies on additive separability. Under multiplicative utility specifications, interaction effects may mechanically arise even in one-dimensional environments. Our contribution is to show that interaction emerges even under additive utility once uncertainty spans orthogonal dimensions.

learning regimes as a function of the information cost, distinguishing regions in which the consumer learns both dimensions, neither dimension, or exactly one dimension. We then study endogenous search order within the intermediate-cost region, where the identity of the single learned dimension is an equilibrium object. Finally, we identify conditions under which equilibrium learning increases ex post variance relative to its ex ante level with strictly positive probability.

To isolate the mechanism, we proceed in two steps. We first analyze a benchmark model in which the consumer's uncertainty lies entirely in the attribute space. We then turn to the main model, in which uncertainty spans both attribute and mapping dimensions. Table 1 summarizes the environments considered.

Environment	ω	α	β	Variance Effect
Attribute-only benchmark	known	uncertain	uncertain	Learning reduces variance
Multidimensional ($\omega \to \beta$)	uncertain	known	uncertain	Variance may increase
Multidimensional ($\beta \to \omega$)	uncertain	known	uncertain	Variance may increase

Table 1: Uncertainty dimensionality and learning effects

5.1 Benchmark: Attribute-level uncertainty only

We begin with a benchmark in which the consumer's uncertainty lies solely in the attribute space. Suppose ω is known, while both α and β are uncertain. Let $\mathrm{Var}(\alpha) = \sigma_\alpha^2 > 0$ and $\mathrm{Var}(\beta) = \sigma_\beta^2 > 0$, with α and β independent. Because utility is linear in the attributes, uncertainty about u arises only from attribute-level variation. Ex ante variance is therefore

$$\mathrm{Var}(u) = \omega^2\sigma_\alpha^2 + (1-\omega)^2\sigma_\beta^2. \tag{10}$$

Learning either attribute eliminates the corresponding variance component. As a result, posterior variance weakly decreases regardless of search order.

Lemma 1. *If the consumer's uncertainty lies solely in the attribute space, learning weakly reduces posterior variance of utility, regardless of search order.*

This benchmark clarifies that non-monotonicity cannot arise in one-dimensional environments. It requires multidimensional uncertainty.

5.2 Main model: Multi-dimensional uncertainty

We now return to the discrete baseline specification introduced in Section 3.2, where uncertainty spans both attribute and mapping dimensions. The salient attribute α is known, $\beta \in \{\mu_\beta \pm \sigma_\beta\}$ with equal probability, and $\omega \in \{0, 1\}$ with $\mathbb{P}(\omega = 1) = q$. Unless otherwise stated, ω and β are independent.

Because utility depends on both the uncertain attribute realization and the uncertain mapping parameter, ex ante dispersion reflects exposure to two orthogonal sources of uncertainty. Using the decomposition derived in Section 4, ex ante variance is

$$\mathrm{Var}(u) = (1-q)\sigma_\beta^2 + q(1-q)(\alpha - \mu_\beta)^2. \tag{11}$$

The first term captures exposure to outcome uncertainty through β, while the second term captures exposure to mapping uncertainty through ω. The interaction between these two components is the source of the non-monotonic learning effects characterized below.

Learning ω first

Consider first the strategy in which the consumer resolves mapping uncertainty before learning about the attribute realization. After paying cost c, she observes whether $\omega = 1$ or $\omega = 0$.

If $\omega = 1$, utility depends only on the known attribute α, and posterior variance collapses to zero. On the other hand, if $\omega = 0$, utility depends entirely on β, so the remaining dispersion equals σ_β^2. Accordingly,

$$\mathrm{Var}(u|\omega) = \begin{cases} 0 & \text{if } \omega = 1, \\ \sigma_\beta^2 & \text{if } \omega = 0. \end{cases} \tag{12}$$

To determine whether learning increases uncertainty, compare this to the ex ante variance in (11). Ex post variance exceeds ex ante variance precisely when $\omega = 0$ and

$$\sigma_\beta^2 > (1-q)\sigma_\beta^2 + q(1-q)(\alpha - \mu_\beta)^2, \tag{13}$$

which simplifies to

$$\sigma_\beta^2 > (1-q)(\alpha - \mu_\beta)^2. \tag{14}$$

Condition (14) highlights the amplification mechanism. If the uncertain attribute is sufficiently volatile relative to the ex ante attribute gap, resolving mapping uncertainty fully exposes outcome uncertainty. In those states, learning reveals that what matters most is precisely the dimension that remains uncertain, thereby increasing perceived dispersion.

Learning β first

Consider next the strategy in which the consumer resolves outcome uncertainty before learning about the mapping parameter. After paying cost c, she observes the realization of β. Conditional on β, the residual uncertainty lies on ω. The product utility equals α if $\omega = 1$ and equals β if $\omega = 0$. Because the dispersion now reflects uncertainty about which of these two values will ultimately determine utility, posterior variance is

$$\mathrm{Var}(u|\beta) = q(1-q)(\alpha - \beta)^2. \tag{15}$$

Comparing (15) to the ex ante variance in (11), ex post variance exceeds ex ante variance when

$$(\alpha - \beta)^2 > \frac{\sigma_\beta^2}{q} + (\alpha - \mu_\beta)^2 \Leftrightarrow \sigma_\beta^2 < q\left(\mu_\beta - \beta\right)\left(2\alpha - \beta - \mu_\beta\right). \tag{16}$$

Condition (16) clarifies the amplification force. If the realized attribute level differs sufficiently from α, then mapping uncertainty becomes more consequential than it was under prior beliefs. In such states, resolving outcome uncertainty increases the economic importance of not knowing which dimension ultimately matters, and overall variance rises.

Equilibrium learning breadth

So far, we have characterized how resolving each dimension affects utility dispersion. We next study when the consumer acquires information at all and, when she does, whether she would resolve one dimension or both. Because learning is costly, equilibrium behavior features three qualitatively distinct regimes as a function of the per-dimension cost c: learning both dimensions, learning neither dimension, and an intermediate region in which the consumer learns exactly one dimension. This intermediate region is where search order becomes an equilibrium object.

We hereafter restrict attention to the more interesting parameter region $\sigma_\beta > |\mu_\beta|$ such that the consumer learns both types of uncertainty when the learning cost diminishes to zero. Furthermore, to sharpen insights, we fix $q = 1/2$. We relax these simplifying assumptions in two separate extensions.

Proposition 1. *Let $\underline{c}$ and $\overline{c}$ be as defined in the proof. In equilibrium,*

1. *if $c \leq \underline{c}$, the consumer acquires both signals in equilibrium (possibly stopping early if she learns $\omega = 1$);*
2. *if $\overline{c} < c$, the consumer acquires no signal;*
3. *if $\underline{c} < c \leq \overline{c}$, the consumer acquires exactly one signal with positive probability.*

Proposition 1 identifies the cost region in which order choice matters: only when the consumer is neither fully uninformed (high c) nor fully informed (low c) can equilibrium learning would lead to resolve exactly one dimension; consequently, the search sequence impacts outcomes. We analyze endogenous ordering for this intermediate cost interval below.

5.3 Endogenous Search Order

We henceforth focus on the intermediate-cost region $c \in (\underline{c}, \overline{c})$ identified in Proposition 1, where the consumer acquires exactly one signal with strictly positive probability. In this region, the identity of the first signal is an equilibrium object. Because learning is sequential, the first signal generates an option value: after observing its realization, the consumer chooses whether to stop and act on her current information or to pay the additional cost c to acquire the second signal. Therefore, she acquires the second signal if and only if its incremental expected value (relative to stopping) weakly exceeds c.

Proposition 2 (Learning prioritization)**.** *Let the thresholds $\underline{\sigma}$ and $\overline{\sigma}$ be as defined in the proof, and consider the intermediate cost range $(\underline{c}, \overline{c})$. The consumer learns (i) only ω if $\sigma_\beta \leq \underline{\sigma}$; and (ii) only β if $\overline{\sigma} < \sigma_\beta$.*

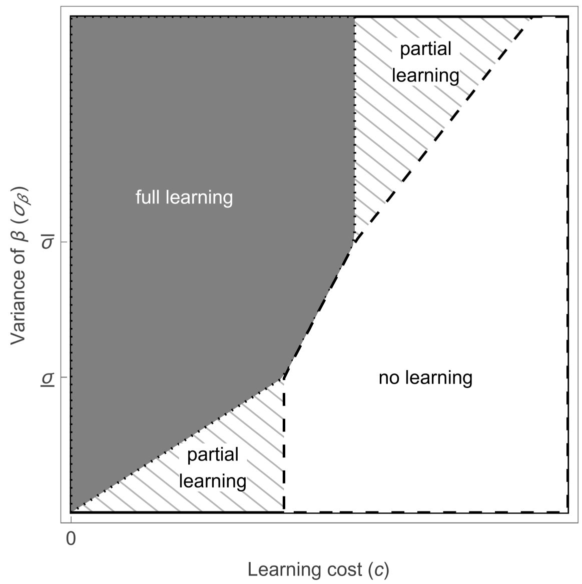

Figure 2: Learning equilibrium; dotted and dashed lines denote $\underline{c}$ and $\overline{c}$, respectively ($\alpha = .4, \mu_\beta = -.3, q = .5$)

Proposition 2 characterizes the conditions under which the consumer optimally chooses to resolve only one dimension of uncertainty (see hashed regions in Figure 2). The choice of uncertainty type to resolve therefore reflects a comparison of option values. Learning ω only is attractive when it prevents wasteful learning about an attribute that may ultimately not matter. Learning β only is attractive when the attribute realization itself is sufficiently informative. In equilibrium, the consumer allocates attention toward the dimension that is more consequential ex ante.

5.4 Equilibrium uncertainty amplification

We now combine the equilibrium learning regime identified in Proposition 1 with the endogenous order choice in Propsoition 2. The key observation is that uncertainty amplification requires the consumer to resolve exactly only one uncertainty dimension. If both ω and β are learned, posterior variance collapses to zero. If neither is learned, posterior variance remains at its ex ante level in (11). Hence, an increase in dispersion can arise only in the intermediate cost regions where equilibrium learning resolves a single dimension, in which case the relevant amplification conditions are those derived in the two learning cases above (see black regions in Figure 3). We delineate the conditions in the following proposition.

Proposition 3 (Equilibrium uncertainty amplification)**.** *In the intermediate cost region, equilib-*

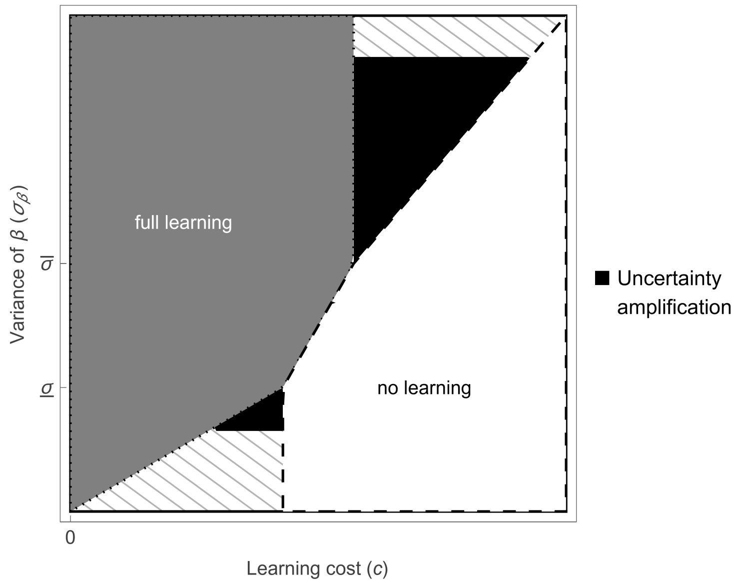

Figure 3: Uncertainty amplification occurs in black region ($\alpha = .4, \mu_\beta = -.3, q = .5$)

rium learning increases perceived utility dispersion even though each learning action is individually informative. Specifically, for any $c \in (\underline{c}, \overline{c})$,

1. *if $|\alpha - \mu_\beta|/\sqrt{2} < \sigma_\beta \leq \underline{\sigma}$, then the equilibrium learning strategy resolves only ω and the ex post variance exceeds ex ante variance with positive probability; and*
2. *if $\overline{\sigma} < \sigma_\beta$, then the equilibrium learning strategy resolves only β, and if this learned $\beta \in \{\mu_\beta \pm \sigma_\beta\}$ satisfies $\sigma_\beta \leq \sqrt{(\mu_\beta - \beta)(2\alpha - \beta - \mu_\beta)/2}$, then the ex post variance exceeds ex ante variance with positive probability.*

The proposition clarifies the conditions when an increase in perceived uncertainty can arise in equilibrium. In the intermediate cost region, the consumer optimally decides to resolve exactly one dimension of uncertainty, which is precisely when amplification can occur. When the consumer learns ω but not β, learning can reveal that the uncertain attribute is the main driver of utility (i.e., the event $\omega = 0$), which makes utility fully sensitive to the residual uncertainty in β. When the consumer learns β but not ω, the learned value of β may turn out to be substantially different from α, which makes it more consequential that the consumer still does not know whether utility will be determined by α or by β. In both cases, learning does not add uncertainty; rather, it makes the remaining uncertainty dimension more salient for utility.

This mechanism provides a fully rational foundation of a Dunning–Kruger type pattern in perceived confidence. The model does not imply that additional information mechanically lowers confidence. Rather, confidence can be non-monotone in learning because partial learning can reveal that the still unresolved dimension of uncertainty is pivotal for utility, increasing posterior dispersion. Consumers who stop early (for example, due to higher information costs) may never uncover this payoff relevance and can therefore remain relatively confident, while consumers who learn more are more likely to reach states in which learning exposes that what remains unknown is precisely what matters most, leading to lower confidence despite more information even under correct Bayesian updating.

Section 6 develops these implications in a more realistic information environment by allowing signals to be noisy and continuous, and by endogenizing the precision of information acquisition. These extensions both sharpen the interpretation and show that the amplification mechanism is robust beyond the discrete baseline.

6 Noisy Continuous Signals and Endogenous Precision

Section 5 established that, under multidimensional uncertainty, equilibrium learning can increase posterior dispersion of utility in the discrete baseline. We now extend the analysis in two ways. First, we replace perfect revelation with noisy continuous signals about β, and show that the amplification logic survives. Second, we allow consumers to choose signal precision endogenously (a rational inattention formulation). Endogenous precision generates systematic differences in how much consumers learn and provides a natural route to cross-sectional patterns in perceived confidence: consumers who acquire more information can become less confident when learning reveals that what remains uncertain is precisely what matters most for utility.

6.1 Noisy Signals and Posterior Utility Dispersion

We focus on a continuous analogue of the baseline environment. The consumer's utility is given by $u = w\,\alpha + (1-w)\beta$, where $w \in \{0, 1\}$ is a random indicator of which attribute matters for utility. As in the discrete model, w captures mapping uncertainty and β captures outcome uncertainty. Let $\mathbb{P}(w = 1) = p$, and assume α is known. The uncertain attribute β follows a normal prior

$$\beta \sim \mathcal{N}(\mu_\beta, \sigma_\beta^2). \tag{17}$$

The consumer can learn about β by observing a noisy signal s, which is

$$s = \beta + \varepsilon, \quad \text{where } \varepsilon \sim \mathcal{N}(0, \lambda^{-1}), \tag{18}$$

where $\lambda \geq 0$ denotes the signal precision (the inverse of noise variance). A higher λ corresponds to a more informative (more precise) signal about β, but it is also more costly to acquire in the endogenous-precision formulation below. The signal is independent of ω, so observing s does not update beliefs about which attribute matters.

Posterior belief updating about β

After observing signal s, the consumer updates her belief about β using Bayesian inference. Since the prior and noise are normal, the posterior for β is also normal. Standard normal updating implies that the posterior variance of β (denoted $\sigma^2_{\beta|s}$) is normal and given:

$$\sigma^2_{\beta|s} \ = \ \frac{\sigma_\beta^2}{1 + \lambda\,\sigma_\beta^2} \ = \left(\lambda + \sigma_\beta^{-2}\right)^{-1} \tag{19}$$

Importantly, this posterior variance $\sigma^2_{\beta|s}$ does not depend on the particular realization of signal s but it only depends on the signal precision λ. Intuitively, a higher λ (more precise signal) leads to a smaller posterior variance of β (more certainty about β). If $\lambda = 0$ (no information), the posterior variance remains σ_β^2 (the prior variance), and as $\lambda \to \infty$ (precise signal), the posterior variance tends to 0 (perfect knowledge of β).

In what follows, as explained in the discrete baseline (Sections 4), we interpret posterior dispersion $\mathrm{Var}(u|s)$ as an operational measure of perceived confidence.

Posterior dispersion of utility after observing s

Even after seeing s, the consumer is still uncertain about her utility because there are two layers of uncertainty: (i) *mapping uncertainty* about which attribute is payoff-relevant (i.e., uncertainty over

w), and (ii) *outcome uncertainty* about the uncertain attribute (i.e., residual uncertainty about β when $w = 0$). Applying the law of total variance yields:

$$\mathrm{Var}(u|s) = (1-p)\,\sigma^2_{\beta|s} + p(1-p)\,(\alpha - \mathbb{E}[\beta|s])^2. \tag{20}$$

The first term $(1-p)\,\sigma^2_{\beta|s}$ is the remaining variance due to imperfect knowledge of β (weighted by the probability $1-p$ that β actually matters). This term *decreases* as λ (signal precision) increases, since more precise information lowers $\sigma^2_{\beta|s}$. The second term $p(1-p)(\alpha - E[\beta|s])^2$ captures the dispersion induced by not knowing whether utility will be α or will instead depend on β. This term can *increase* with precision because a more precise signal can move $\mathbb{E}[\beta|s]$ farther away from α in some realizations, thereby widening the payoff difference between the $w = 1$ and $w = 0$ states. Thus, as in the discrete model, learning about the outcome dimension can increase the payoff relevance of mapping uncertainty. Importantly, a higher λ reduces residual uncertainty about β through $\sigma^2_{\beta|s}$, but it also makes the posterior mean $\mathbb{E}[\beta \mid s]$ respond more strongly to the realized signal s. As a result, for some realizations s, greater precision can widen the implied gap between the two candidate utility outcomes (utility equals α when $w = 1$ and depends on β when $w = 0$), increasing posterior dispersion of u even though the signal about β becomes more accurate. We summarize this state-dependent effect of precision in the following proposition.

Proposition 4 (State-dependent amplification under noisy learning)**.** *Given any $\lambda > 0$, there exist signal realizations s such that increasing precision λ raises posterior dispersion* $\mathrm{Var}(u \mid s)$*.*

Proposition 4 formalizes the same intuition as in the discrete baseline. More precise learning about the outcome dimension can, for some realizations, increase the payoff relevance of mapping uncertainty by widening the gap between the two candidate utility outcomes.

6.2 Endogenous Precision Choice

We now endogenize the informativeness of the signal by allowing the consumer to choose precision λ at a cost. We interpret λ as an attention intensity: higher precision corresponds to more attentive processing of information about β but incurs a higher attention cost. A key modeling issue is that, under risk neutrality, the consumer does not value variance reduction per se; she values information because it improves subsequent decisions. We therefore adopt a simple two-stage microfoundation in which precision affects expected utility through its effect on a downstream choice.

We solve the consumer's problem via backward induction. After observing s, the consumer chooses between two actions. Action A yields a sure payoff α, while action B yields payoff β. The consumer chooses B if $\mathbb{E}[\beta \mid s] \geq \alpha$ and otherwise selects A. This reduced-form choice stage captures the standard idea that, even for a risk-neutral agent, more precise information improves decision quality by sharpening posterior beliefs.

Anticipating the Stage-2 decision rule, the consumer's Stage-1 problem is

$$\max_{\lambda \geq 0} \mathbb{E}_{s(\lambda)} \left[\max\{\alpha, \mathbb{E}[\beta \,|\, s]\} \right] - C(\lambda). \tag{21}$$

This formulation delivers an endogenous precision choice without assuming that the consumer directly minimizes variance.

Note that the precision choice of λ in (21) is pinned down by expected utility, not by taking variance reduction as the objective. Nevertheless, λ is closely linked to posterior dispersion in our setting. Because λ governs signal informativeness, it directly shapes the posterior terms in (20), which we use as an operational measure of perceived confidence. In particular, as summarized in Proposition 4, increases in precision can raise posterior dispersion for some signal realizations by amplifying the payoff relevance of the unresolved mapping uncertainty. With a convex attention cost $C(\lambda)$, this implies that the optimal precision can be interior, $\lambda > 0$ and finite. This interior precision choice naturally generates partial learning, and in extensions it provides a simple way to connect differences in information acquisition (e.g., due to different attention costs or processing efficiency) to cross-sectional patterns in perceived confidence.

6.3 Expected Posterior Dispersion and Comparative Statics

Although λ is chosen to maximize expected utility in (21), it is useful to characterize how our variance-based confidence measure varies with precision. In particular, we compute the *ex ante* expected posterior dispersion $\mathbb{E}[\mathrm{Var}(u \mid s)]$ implied by a given precision λ. This object is not the consumer's objective; it summarizes how information processing translates into perceived confidence on average. To compute $\mathbb{E}[\mathrm{Var}(u \mid s)]$, we start from the posterior-variance decomposition $\mathrm{Var}(u \mid s) = (1-q)\sigma^2_{\beta|s} + q(1-q)(\alpha - \mathbb{E}[\beta \mid s])^2$ from Equation (20), and take expectations over the signal realization s. The first term is straightforward because $\sigma^2_{\beta|s}$ depends only on λ and not on s. The second term depends on s through the posterior mean $\mathbb{E}[\beta \mid s]$, so we compute the key term $\mathbb{E}_s[(\alpha - \mathbb{E}[\beta \mid s])^2]$, which captures the ex ante expected squared gap between the two candidate utility outcomes implied by mapping uncertainty.

Using (19) under normal updating and imposing $\mu_\beta = \alpha$ for simplicity,[5]

$$\mathbb{E}_s\left[(\alpha - \mathbb{E}[\beta \mid s])^2\right] = \sigma_\beta^4 \frac{\lambda}{1+\lambda\sigma_\beta^2}, \tag{22}$$

Substituting into (20) yields

$$\mathbb{E}\left[\mathrm{Var}(u \mid s)\right] = \frac{(1-q)\sigma_\beta^2(1+q\lambda\sigma_\beta^2)}{1+\lambda\sigma_\beta^2}. \tag{23}$$

Equation (23) shows that expected posterior dispersion decreases with precision λ, even though dispersion can increase with λ for some signal realizations (as captured by Proposition 4).

Proposition 5 (Expected dispersion under precision λ)**.** *Expected posterior dispersion* $\mathbb{E}[\mathrm{Var}(u \mid s)]$ *is decreasing in λ for fixed (q, σ_β^2). However, posterior dispersion* $\mathrm{Var}(u \mid s)$ *can increase in λ for some realizations s (Proposition 4).*

Proposition 5 highlights a useful distinction: while greater precision reduces dispersion on average, it can increase posterior dispersion for some realizations by amplifying the payoff relevance of mapping uncertainty. This state-dependent channel shapes the marginal value of additional precision and therefore the consumer's optimal attention choice.

Optimal attention choice

We now characterize the consumer's optimal attention intensity, captured by the signal precision λ. The cost of attention is assumed to be a quadratic $C(\lambda) = \kappa \cdot \lambda^2$ with $\kappa > 0$. Recall the consumer chooses λ to maximize the expected utility in (21). Define the gross benefit from signal precision

[5] Allowing $\mu_\beta \neq \alpha$ adds a constant term to $\mathbb{E}_s[(\alpha - \mathbb{E}[\beta \mid s])^2]$ and does not change the qualitative results.

λ as

$$G(\lambda) \equiv \mathbb{E}_{s(\lambda)}\left[\max\{\alpha, \mathbb{E}[\beta \mid s]\}\right]. \tag{24}$$

Then, the consumer's problem can be rewritten as

$$\max_{\lambda \geq 0} \; G(\lambda) - \kappa\lambda^2. \tag{25}$$

An interior optimum $\lambda^* > 0$ satisfies the first-order condition

$$G'(\lambda^*) = 2\kappa\,\lambda^*. \tag{26}$$

which equates the marginal decision-improvement benefit of precision to the marginal attention cost. Because the attention cost is convex while the informational benefit exhibits diminishing returns, the optimal precision can be interior: $\lambda^* > 0$ and finite.

Proposition 6 (Existence and comparative statics of optimal precision)**.** *For $\kappa > 0$, there exists an interior optimal precision $\lambda^* > 0$*

$$\frac{(1-q)^2\,\sigma_\beta^4}{(1+\sigma_\beta^2\lambda)^2} = 2\kappa\,\lambda. \tag{27}$$

Moreover, λ^* increases with σ_β^2 and decreases with κ.

7 Conclusion

This paper investigates a counterintuitive possibility in consumer learning wherein acquiring more information can amplify perceived uncertainty about a product's utility. We develop a Bayesian model of multidimensional evaluation where consumers optimize the sequence of uncertainties to resolve. Unlike standard one-dimensional learning models in which additional information monotonically reduces posterior dispersion, our analysis shows that when uncertainty spans both outcome and mapping dimensions, learning can reallocate uncertainty toward what matters most for utility. As a result, consumer's confidence may fall with information.

Moreover, we show that this amplification mechanism is robust beyond the discrete baseline:

under noisy signals and endogenous precision choice, higher signal precision can increase posterior dispersion for some realizations, and optimal attention choice can generate systematic differences in confidence across consumers. Taken together, these results provide a rational micro-foundation for Dunning-Kruger-type confidence patterns without invoking biased updating or other behavioral assumptions.

The analysis is subject to several limitations. To isolate the amplification mechanism, the baseline model abstracts from pricing and competition. While the continuous and endogenous-precision extensions demonstrate robustness, the model remains stylized in its treatment of attribute space and market structure. While these simplifications are deliberate, they leave open important questions about how the mechanism interacts with potentially more realistic institutional settings. Future research can extend this framework in several directions. Embedding multidimensional uncertainty into competitive models would allow firms' pricing and disclosure strategies to be endogenized and linked to learning dynamics. Empirical work using experiments or field data could test whether partial learning increases perceived uncertainty and delays purchase when mapping uncertainty is salient.

Declarations

Funding and competing interests: All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript. The authors have no funding to report.

References

Branco, F., Sun, M., and Villas-Boas, J. M. (2016). Too much information? information provision and search costs. *Marketing Science*, 35(4):605–618.

Caplin, A. and Dean, M. (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7):2183–2203.

Carmon, Z., Wertenbroch, K., and Zeelenberg, M. (2003). Option attachment: When deliberating makes choosing feel like losing. *Journal of Consumer Research*, 30(1):15–29.

DeGroot, M. H. (1970). *Optimal Statistical Decisions*. McGraw-Hill, New York.

Diehl, K. and Poynor, C. (2004). Great expectations?! assortment size, expectations, and satisfaction. *Journal of Marketing Research*, 41(3):312–322.

Dukes, A. and Liu, L. (2016). Online shopping intermediaries and the strategic design of search environments. *Management Science*, 62(4):1064–1077.

Dzyabura, D. and Hauser, J. R. (2019). Recommending products when consumers learn their preference weights. *Marketing Science*, 38(3):417–441.

Erdem, T. and Keane, M. P. (1996a). Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Science*, 15(1):1–20.

Erdem, T. and Keane, M. P. (1996b). Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Science*, 15(1):1–20.

Guo, L. and Zhang, J. (2012). Consumer deliberation and product line design. *Marketing Science*, 31(6):995–1007.

Hauser, J. R. and Wernerfelt, B. (1990). An evaluation cost model of consideration sets. *Journal of Consumer Research*, 16(4):393–408.

Kruger, J. and Dunning, D. (1999). Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology*, 77(6):1121.

Kuksov, D. and Villas-Boas, J. M. (2010). When more alternatives lead to less choice. *Marketing Science*, 29(3):507–524.

Lee, J.-Y., Shin, J., and Yu, J. (2025). Communicating attribute importance under competition. *Marketing Science*.

Lee, L. and Ariely, D. (2011). Searching for confidence: The impact of comparative information on consumer confidence. *Journal of Marketing Research*, 48(4):686–701.

Li, J. and Hitt, L. (2021). Attention and product design. *Marketing Science*, 40(2):209–228.

Li, X., Li, Y., and Shi, M. (2019). Managing consumer deliberations in a decentralized distribution channel. *Marketing Science*, 38(1):170–190.

Liu, Q. and Dukes, A. (2022). Information provision and consumer search. *Marketing Science*, 41(3):405–424.

Matejka, F. and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–298.

Moore, D. A. and Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, 115(2):502–517.

Puntoni, S., De Langhe, B., and Van Osselaer, S. (2009). Bilingualism and the emotional intensity of advertising language. *Journal of Consumer Research*, 36(6):1012–1025.

Shugan, S. M. (1980). The cost of thinking. *Journal of consumer Research*, 7(2):99–111.

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

Sriram, S., Chintagunta, P., and Neelamegham, R. (2006). Effects of brand preference, product attributes, and marketing mix variables in technology product markets. *Marketing Science*, 25(5):440–456.

Sriram, S. and Kalwani, M. U. (2007). Optimal advertising and promotion budgets in dynamic markets with brand equity as a mediating variable. *Management Science*, 53(1):46–60.

Tormala, Z. L. and Petty, R. E. (2006). What doesn't kill me makes me stronger: The effects of resisting persuasion on attitude certainty. *Journal of Personality and Social Psychology*, 91(6):1009–1024.

Appendix

A1 Proofs

A1.1 Proof of Lemma 1

The proof is trivial and omitted.

Proof of Proposition 1

The consumer's expected payoff for each of the strategies is as follows:

- If the consumer learns $\omega = 1$, then her expected utility is α^+, because she chooses the greater of α or 0. If she learns $\omega = 0$, then she decides whether to learn β. If she does not learn β, her expected utility is

$$\frac{1}{2}\mu_\beta^+ \tag{A1}$$

and if she does learn, it is

$$\frac{1}{2}(\mu_\beta - \sigma_\beta)^+ + \frac{1}{2}(\mu_\beta + \sigma_\beta)^+ - c. \tag{A2}$$

Overall, if she learns ω first, her net expected payoff is

$$u(\omega\text{-first}) = q\alpha^+ + (1-q)\max\left\{\mu_\beta^+, \frac{1}{2}(\mu_\beta - \sigma_\beta)^+ + \frac{1}{2}(\mu_\beta + \sigma_\beta)^+ - c\right\} - c. \tag{A3}$$

- Similarly, if she learns β first, her net expected payoff is

$$u(\beta\text{-first}) = \left(\sum_{\beta \in \{\mu_\beta \pm \sigma_\beta\}} \max\left\{(q\alpha + (1-q)\beta)^+, q\alpha^+ + (1-q)\beta^+ - c\right\}/2\right) - c. \tag{A4}$$

- Finally, if she does not learn anything, her expected utility is

$$u(\emptyset) = (q\alpha + (1-q)\mu_\beta)^+. \tag{A5}$$

Consumer learns both dimensions if and only if (i) consumer conducts first search (i.e., $\max\{u(\omega\text{-first}), u(\beta\text{-first})\} > u(\emptyset)$), and (ii) conducts second search; (i.e., if $u(\omega\text{-first}) > u(\beta\text{-first})$, then $\mu_\beta^+ < \frac{1}{2}(\mu_\beta - \sigma_\beta)^+ + \frac{1}{2}(\mu_\beta + \sigma_\beta)^+ - c$; otherwise, $(q\alpha + (1-q)\beta)^+ < q\alpha^+ + (1-q)\beta^+ - c$ for $\beta \in \{\mu_\beta \pm \sigma_\beta\}$). All conditions reduce to c being small, so defining the minimum of all the cost thresholds that satisfy the inequalities above as $\underline{c}$, we obtain that the consumer learns both dimensions if and only if $c < \underline{c}$.

Similarly, the consumer learns neither dimensions if and only if $\max\{u(\omega\text{-first}), u(\beta\text{-first})\} < u(\emptyset)$. The left-hand side is decreasing in c whereas the right-hand side is independent of c. Furthermore, the left-hand side is greater at $c = 0$ and the right-hand side is greater for $c \uparrow \infty$; therefore, there exists a unique threshold $\bar{c}$ such that $\max\{u(\omega\text{-first}), u(\beta\text{-first})\} < u(\emptyset)$ if and only if $c > \bar{c}$.

By construction, for any $c \in (\underline{c}, \bar{c})$, the consumer only partially resolves uncertainty. This completes the proof.

Proof of Proposition 2

First, we can rule out the following degenerate parameter regions. If $\mu_\beta + \sigma_\beta < 0$ and $\alpha < -$, then the product utility can never exceed 0. Therefore, the consumer chooses outside option without learning. Similarly, if $\mu_\beta - \sigma_\beta > 0$ and $\alpha > 0$, then the product utility always exceeds 0. Therefore, the consumer chooses the product without learning.

If either (i) $\mu_\beta + \sigma_\beta < 0$ and $\alpha > 0$, or (ii) $\mu_\beta - \sigma_\beta > 0$ and $\alpha < 0$, then in the intermediate c region where consumer learns only one dimension, she always learns ω first.

Next, we focus on $\mu_\beta - \sigma_\beta < 0 \le \mu_\beta + \sigma_\beta$. Then in the intermediate c region where consumer learns only one dimension, the consumer's expected payoff from only learning ω is

$$u(\omega\text{-only}) = q\alpha^+ + (1-q)\mu_\beta^+ - c, \tag{A6}$$

and her the expected payoff from only learning β is

$$u(\beta\text{-only}) = -c + \sum_{\beta \in \{\mu_\beta \pm \sigma_\beta\}} (q\alpha + (1-q)\beta)^+/2. \tag{A7}$$

Due to convexity of the $(\cdot)^+$ operator, $u(\beta\text{-only})$ is strictly increasing in σ_β, whereas $u(\omega\text{-only})$ is independent of σ_β. Moreover, as $\sigma_\beta \downarrow 0$, $u(\beta\text{-only})$ simplifies to $(q\alpha+(1-q)\mu_\beta)^+-c$, which is smaller than $q\alpha^+ + (1-q)\mu_\beta^+ - c$. Therefore, by the Intermediate Value Theorem and strict monotonicity, there exists a unique threshold $\tilde{\sigma}_\beta$ such that learning β-only yields higher payoff than learning ω-only if and only if $\sigma_\beta > \tilde{\sigma}_\beta$. Let $\sigma_\beta^i = \inf\{\sigma_\beta > 0 : \underline{c} = \overline{c}\}$ and $\sigma_\beta^s = \sup\{\sigma_\beta > 0 : \underline{c} = \overline{c}\}$. Defining $\underline{\sigma} = \min\left\{\tilde{\sigma}_\beta, \sigma_\beta^i\right\}$ and $\overline{\sigma}_\beta = \max\left\{\tilde{\sigma}_\beta, \sigma_\beta^s\right\}$.

With $q = 1/2$, the thresholds simplify as follows:

$$\underline{\sigma} = \begin{cases} |\alpha - \mu_\beta| & \text{if } -\mu_\beta < \alpha < 0 \text{ or } 0 < \alpha < -\mu_\beta, \\ 2|\mu_\beta| & \text{otherwise} \end{cases}$$

$$\overline{\sigma} = \begin{cases} |\alpha - \mu_\beta| & \text{if } -\mu_\beta < \alpha < 0 \text{ or } 0 < \alpha < -\mu_\beta, \\ |3\alpha + \mu_\beta| & \text{otherwise}, \end{cases}$$

$$\underline{c} = \begin{cases} \frac{1}{2}((\mu_\beta - \sigma_\beta)^+ + (\mu_\beta + \sigma_\beta)^+) - \mu_\beta^+ & \text{if } \sigma_\beta \le \underline{\sigma} \\ \frac{1}{6}\left(-2(\alpha + \mu_\beta)^+ + 2\alpha^+ + \mu_\beta + \sigma_\beta\right) & \text{if } \underline{\sigma} < \sigma_\beta \le \overline{\sigma}, \\ \max\left(\frac{1}{2}\left(\alpha^+ - (a + \mu_\beta - \sigma_\beta)^+\right), \frac{1}{2}\left(-(\alpha + \mu_\beta + \sigma_\beta)^+ + \alpha^+ + \mu_\beta + \sigma_\beta\right)\right) & \text{otherwise} \end{cases}$$

$$\overline{c} = \begin{cases} \frac{1}{2}\left(\alpha^+ - (\alpha + \mu_\beta)^+ + (\mu_\beta)^+\right) & \text{if } \sigma_\beta \le \underline{\sigma}, \\ \underline{c} & \text{if } \underline{\sigma} < \sigma_\beta \le \overline{\sigma}, \\ \frac{1}{2}\left(\frac{1}{2}(\alpha + \mu_\beta - \sigma_\beta)^+ + \frac{1}{2}(\alpha + \mu_\beta + \sigma_\beta)^+ - (\alpha + \mu_\beta)^+\right) & \text{otherwise.} \end{cases}$$

A1.2 Proof of Proposition 3

From Proposition 2, ex post variance cannot be higher than ex ante variance if either $c \le \underline{c}$ (because full learning reduces ex post variance to zero) or $\overline{c} < c$ (because variance remains unchanged without learning).

However, under partial learning for intermediate $c \in (\underline{c}, \overline{c})$, conjunction with Conditions (14) and (16) yields the necessary and sufficient condition for the uncertainty amplification effect.