Violations of Bayesian updating within Expected Utility

Elias Tsakas

Maastricht University

Vassili Vergopoulos

Université Paris-Panthéon-Assas

November, 2026

Preliminary draft

Latest version of the paper.

Abstract

We study violations of Bayesian updating within a standard setting of preferences that satisfy the Subjective Expected Utility axioms. We show that Bayesian updating will be corroborated if and only if utility function is signal-independent. In this sense, we characterize when violations of Bayesianism can be attributed to belief elicitation errors due to failure to take into account that people may care about the information they receive or the associated beliefs in an intrinsic way, like in the extensive literature on belief-dependent motives and emotionally-dependent beliefs.

Keywords: Subjective Expected Utility, Bayesian updating, signal-dependent utility, belief misspecification, dynamic consistency.

1. Introduction

It is widely accepted within the decision-theoretic literature that Subjective Expected Utility (SEU) preferences always guarantee Bayesian updating (Karni and Schmeidler, 1991; Epstein and Le Breton, 1993; Ghirardato, 2002). In fact, this is true even if we relax the assumption of utilities being state-independent (Karni *et al.*, 1983; Lu, 2019; Tsakas, 2026). However, what decision-theorists have failed to take into account is the possibility of utilities being signal-dependent. This is likely to occur is situations where the agent cares about the type of news she will receive, or has belief-dependent utilities which are typically triggered by different emotional responses (Caplin and Leahy, 2001; Köszegi, 2006; Köszegi and Rabin, 2009; Golman *et al.*, 2022).

In this paper, we will allow for belief-dependent utilities within the standard SEU framework. This will result into a class of SEU representations, which we will call *signal-dependent SEU*, analogously to state-dependent SEU representations. Let us illustrate such a representation with a simple example.

Consider a (female) voter with SEU preferences and prior beliefs that assign equal probability to the two states, i.e., to Republicans winning (D) and to Republicans winning the election (R). The voter receives information from a forecaster who either predicts a Democratic win (d) or a Republican win (r). What is objectively known is that the forecaster correctly predicts the winner of elections 75% of the times. The voter is indifferent between the two candidates, but she despises the forecaster who used to be her husband before their ugly divorce, and therefore she receives higher utility if the forecaster is wrong. This is all reflected below: on the left-hand side table we depict her vNM utility for monetary payoffs q), while on the right-hand side we depict her joint beliefs, which are consistent with the objective likelihoods of the forecaster being correct.

	d	r
D	q	$3q$
R	$3q$	q

Utility from q

	d	r
D	$3/8$	$1/8$
R	$1/8$	$3/8$

Joint belief

Then, we ask the following question: if we elicit the voter's prior and posterior beliefs using a standard elicitation mechanism, will Bayesian updating be satisfied? Our main premise is that the belief elicitation mechanism can only condition payments on the state-realization, but not on the signal realization (Schlag and van der Weele, 2013; Schotter and Trevino, 2014).

This is consistent with both theory and practice, where SEU beliefs are identified with the use of traditional choice data, i.e., though observed choices among acts over the state space.

What is well-known is that all standard elicitation mechanisms normalize the beliefs, by taking a state-independent SEU representation. Notice that in our previous example, when eliciting the prior, utilities are already state-independent as the utility functions are weighted by the likelihood of the signals, meaning that both conditional on D and conditional on R the utility from receiving q is equal to $3q/4$. Hence, the elicited prior will assign probability 50% to each of teh states. However, due to signal-dependence, when we condition on d, the utility function is no longer state-independent, i.e., the utility at D will be q, whereas at R it will be $3q$. This means that the posterior belief given d will again assign probability 50% to each of the states. This leads us to the (false) conclusion that Bayesian updating is violated.

In our main result, we show that this is not an artefact of this example. In fact, Bayesian updating will be corroborated if and only if utilities are signal-independent. And while sufficiency is not surprising, necessity is far from obvious. A direct implication of the previous result is that —as long as we trust that the agent understands the information structure— signal independence of the utility function can be tested with traditional choice data. In this sense, we can identify within the SEU framework when Bayesian updating is violated.

Our work relates to several different literatures. Obviously, it is part of the literature that connects Bayesian updating, dynamic consistency and consequentialism (Karni and Schmeidler, 1991; Epstein and Le Breton, 1993; Ghirardato, 2002), in that it provides foundations for Bayesian updating while maintaining all the power of SEU. We also show how the violations of Bayesian updating relate to the large literature of belief-dependent utilities (Caplin and Leahy, 2001; Köszegi, 2006; Köszegi and Rabin, 2009; Golman *et al.*, 2022). Conceptually, we also provide a clear distinction between state-dependent preferences (where people care about ground truth inherently) and signal-dependent preferences (where they care about the information they receive). This provides a direct link between the literature on belief-dependent utilities and state-dependent utilities (Karni *et al.*, 1983; Lu, 2019; Tsakas, 2026). Furthermore, we contribute to the literature on belief updating biases, by providing a test that identifies the false negatives (Benjamin, 2019). But perhaps the literature which is closest to our paper is the one on misspecification, which also attempts to rationalize belief updating which appears to violate Bayesianism (Shmaya and Yariv, 2016; Bohren and Hauser, 2025, 2026; Molavi, 2026). The key difference to our work is that they assume both state-independence and signal-independence in the utility functions and they permit for mistakes in the interpretation of the objectively given likelihoods, whereas we assume that likelihoods are correctly interpreted and we permit belief measurement errors due to signal-independence. In this sense, conceptually, we attribute (false) rejections of Bayesianism to very different cognitive mechanisms.

2. World-dependent Subjective Expected Utility

Consider a finite set of states of nature Θ, and a stochastic signal $\pi : \Theta \to \Delta(S)$ where S is the finite set of signal realizations such that $|S| \leq |\Theta|$. The signal is assumed to be full-support, i.e., $\pi(s|\theta) > 0$ for all $\theta \in \Theta$ and all $s \in S$. Moreover, we assume that the $|\Theta| \times |S|$ matrix Π with entries $\pi(s|\theta)$ is full-rank, i.e., we have $\text{rank}(\Pi) = |S|$. Throughout the paper, we define the set of states of the world by

$$\Omega := \Theta \times S.$$

In order to avoid confusion, we will henceforth call each $\theta \in \Theta$ a state and each $\omega \in \Omega$ a world.

Definition 1. A (full-support) belief $\mu \in \Delta(\Omega)$ is said to be consistent with π if

$$\mu(s|\theta) = \pi(s|\theta) \tag{1}$$

for every world $(\theta, s) \in \Omega$. The set of beliefs that are consistent with π is denoted by $\mathcal{M}_\pi$. $\triangleleft$

Take a convex and closed set of consequences $Q \subseteq \mathbb{R}^n$. The most common special case is the AA framework, where $Q = \Delta(X)$ for some finite set of outcomes X (Anscombe and Aumann, 1963). An act is a mapping

$$f : \Omega \to Q$$

We denote by $f_\omega \in Q$ the consequence that act f yields at world ω. This means that theoretically it is possible to condition the outcome on both Θ and S. Nonetheless, in practice, ourselves (viz., the analyst) will only be able to observe data of choices among Θ-measurable acts, i.e., acts in

$$\mathcal{F}_\Theta := \big\{ f \in \mathcal{F} : f_{\theta,s_1} = f_{\theta,s_2} \text{ for all } s_1, s_2 \in S \big\}.$$

Convex combinations of acts are defined in the usual way, i.e., $(\lambda f + (1-\lambda)g)_\omega = \lambda f_\omega + (1-\lambda) g_\omega$. For any event $E \subseteq \Omega$, we denote by $f_E g \in \mathcal{F}$ the act that agrees with f at every world in E and with g at every world in the complement $\neg E$.

Take a (female) agent with preferences $\succeq$ over $\mathcal{F}$. As usual, the strict preference relation $\succ$ and the indifference relation $\sim$ are respectively given by the asymmetric and symmetric part of $\succeq$. For an event $E \subseteq \Omega$ we define the conditional preferences $\succeq_E$ in the usual way: $f \succeq_E g$ if and only if $f_E h \succeq g_E h$ for every $h \in \mathcal{F}$. For notation simplicity, we denote by $\succeq_\omega := \succeq_{\{\omega\}}$ the preferences (over Q) at world ω.

The preference relation $\succeq$ is assumed to satisfy the basic axioms of Anscombe and Aumann (1963), with a strengthening of the non-triviality axiom which guarantees that all worlds are relevant:

(A_0) Full-support — For every $\omega \in \Omega$: There are $f, g, h \in \mathcal{F}$ such that $f_{\{\omega\}} h \succ_\Theta g_{\{\omega\}} h$

(A_1) Completeness — For every $f, g \in \mathcal{F}$: $f \succeq g$ or $g \succeq f$

(A_2) Transitivity — For every $f, g, h \in \mathcal{F}$: if $f \succeq g$ and $g \succeq h$, then $f \succeq h$

(A_3) Continuity — For every $f, g, h \in \mathcal{F}$: if $f \succ g \succ h$, then $g \sim \lambda f + (1-\lambda)h$ for some $\lambda \in (0,1)$

(A_4) Independence — For every $f, g, h \in \mathcal{F}$ and every $\lambda \in (0,1)$: $f \succ g$ if and only if $\lambda f + (1-\lambda)h \succeq \lambda g + (1-\lambda)h$

(A_5) State-independence — For every $\omega, \omega' \in \Omega$: $\succeq_\omega = \succeq_{\omega'}$

It is well-known that $(A_0)-(A_5)$ guarantee that there is a world-independent utility function $\bar{u} : Q \to \mathbb{R}$ and a full-support belief $\bar{\mu} \in \Delta(\Omega)$ such that $\succeq$ are represented by

$$\mathbb{E}_{\bar{\mu}}\big(\bar{u}(f)\big) = \sum_{\omega \in \Omega} \bar{\mu}(\omega)\bar{u}(f_\omega), \tag{2}$$

i.e., we will have $f \succeq g$ if and only if $\mathbb{E}_{\bar{\mu}}(\bar{u}(f)) \geq \mathbb{E}_{\bar{\mu}}\big(u(g)\big)$. The world-independent SEU representation $(\bar{u}, \bar{\mu})$ is unique up to positive affine transformation of $\bar{u}$.

As it is well-known $(\bar{u}, \bar{\mu})$ is not the only SEU representation. In fact, for every full-support belief, there exists some world-dependent utility function so that the two together represent $\succeq$. In particular, for any full-support belief μ, define the world-dependent utility function

$$u_\omega := \beta_\omega \cdot \bar{u}, \tag{3}$$

where $\beta_\omega := \bar{\mu}(\omega)/\mu(\omega) > 0$ is the world multiplier parameter. Then, notice that the pair (u, μ) is another SEU representation of $\succeq$. This is because, for each $f \in \mathcal{F}$ we have:

$$\mathbb{E}_\mu(u(f)) = \sum_{\omega \in \Omega} \mu(\omega)u_\omega(f_\omega) = \sum_{\omega \in \Omega} \bar{\mu}(\omega)\bar{u}(f_\omega) = \mathbb{E}_{\bar{\mu}}(\bar{u}(f)).$$

Once we pin down the actual belief, u is uniquely identified up to positive linear transformations (Tsakas, 2026). Hence, the set of SEU representation is isopmorphic to the set of full-support beliefs.

The latter implies that our earlier condition regarding the agent having a correct model in mind about the information structure can be rewritten as a restriction in the set of SEU representations.

Definition 2. A SEU representation (u, μ) is consistent with π if $\mu \in \mathcal{M}_\pi$. Denote the set of such representations by $\mathbb{SEU}_\pi$. ◁

Notice that there is always a SEU representation consistent with π. To see this, fix an arbitrary marginal μ_Θ and construct the joint belief $\mu(\theta, s) := \mu_\Theta(\theta)\pi(s|\theta)$. Then, find the unique utility function which together with μ represents $\succeq$.

The fact that utilities are uniquely identified (once we have pinned down an arbitrary set of beliefs) means that we can formulate conditions on the tastes of the agent. Particularly interesting conditions about the dependence of the utility function on the two dimensions of the world. State-independence has been extensively studied in the literature (Karni *et al.*, 1983; Lu, 2019; Tsakas, 2026), and corresponds to utility functions that are S-measurable. What we introduce in this paper is signal-independence.

Definition 3. A SEU representation (u, μ) is signal-independent if u is Θ-measurable, i.e.,

$$u_{\theta,s_1} = u_{\theta,s_2} \tag{4}$$

for any $\theta \in \Theta$ and any two $s_1, s_2 \in S$. Denote the set of such representations by $\mathbb{SEU}_\Theta$. $\triangleleft$

Examples of signal-dependent utilities include instances where the agent cares about the information she receives. For instance, suppose that she dislikes hearing bad news regardless of whether her tastes are state-dependent or not. Or she may attach intrinsic value to hearing news that reflect the truth. Or even, she may want the information source to be discredited. Notice that we can explicitly disentangle motives driven by tastes about the state space from motives which are driven from tastes about the information.

Before proceeding, we should clarify one important point about the nature of beliefs. Throughout the literature, there are two competing views. According to the traditional decision-theoretic approach, beliefs are simply a parameter within the SEU representation (Grant, 2008; Karni, 2014), and therefore the concept of actual beliefs is meaningless in the first place. More recently, a second view has emerged, according to which beliefs are a primitive that describes the agent's quantification of uncertainty, meaning that it is perfectly fine to talk about the actual belief. This approach has been driven by advances in behavioral economics (Tsakas, 2026) and has also been adopted in the recently surging literature on misspecification (Bohren and Hauser, 2025). For the purposes of this paper both approaches can be taken without affecting the main message. From now onwards, we will implicitly adopt to the second one, thus often referring to the actual belief and the corresponding actual utility, and towards the end we will go back to revisit the classical approach (Section 4.2).

3. Bayesian updating and signal-independent utilities

Throughout the literature, beliefs about Θ are elicited with the use of mechanisms, which incentivize truth-telling by means rewarding accurate reports and punishing inaccurate ones (see

Section A). Notably, all standard mechanisms rely on observing choices among Θ-measurable acts, and elicit the marginal of the state-independent belief,

$$\bar{\mu}_\Theta := \text{marg}_\Theta \, \bar{\mu}.$$

The same is true for posterior beliefs, i.e., all standard mechanisms will elicit $\bar{\mu}(\cdot|s)$ for each $s \in S$. And this brings us to the main question of the paper: when will it be the case that our Expected Utility maximizer appears to violate Bayesian updating?

Definition 4. We say that Bayesian updating is corroborated if, for all $(\theta, s) \in \Omega$,

$$\bar{\mu}(\theta|s) = \frac{\bar{\mu}_\Theta(\theta)\pi(s|\theta)}{\sum_{\theta' \in \Theta} \bar{\mu}_\Theta(\theta')\pi(s|\theta')}.$$

Otherwise, Bayesian updating is violated. ◁

So, if we maintain the approach that an actual belief exists, corroborating the Bayesian updating does not mean that we have elicited the prior and the posterior beliefs correctly. It may very well be the case that we have measured both of them wrongly, and yet the errors cancel out each other so that the agent is still classified as Bayesian. But on the flip side, it may also be the case that Bayesian updating is wrongly rejected due to the fact that the elicitation errors of the prior and the posterior do not cancel out each other. The following result characterizes when Bayesian updating is corroborated and when it is violated.

Theorem 1. *Let $\succeq$ satisfy the AA axioms $(A_0)-(A_6)$. Then, Bayesian updating is corroborated if and only if the utility function is signal-independent. Formally, the following hold:*

(a) Bayesian updating is corroborated if and only if $\mathbb{SEU}_\pi = \mathbb{SEU}_\Theta$.

(b) Bayesian updating is rejected if and only if $\mathbb{SEU}_\pi \cap \mathbb{SEU}_\Theta = \emptyset$.

The theorem relies on an intermediate result, according to which either $\mathbb{SEU}_\pi \cap \mathbb{SEU}_\Theta = \emptyset$ or $\mathbb{SEU}_\pi = \mathbb{SEU}_\Theta$ (Lemma B3). Overall, the previous result shows that, contrary to what has been thought so far, Bayesian updating is not dual to SEU alone, but rather to signal-independent SEU. Most importantly, signal independence is a characteristic of $\succeq$, meaning that it can be tested with traditional choice data, i.e., choices over Θ-measurable acts. Hence, Bayesian updating becomes a testable hypothesis within SEU which can be falsified with traditional choice data.

Of course, if we somehow had access to choice data from choices among all acts, we would be able to elicit the likelihood function, i.e. the conditional beliefs about S given Θ. In this case, any standard mechanism would yield $\bar{\mu}(\cdot|\theta)$. Then, it is not difficult to verify that Bayesian updating would be corroborated if and only if $\bar{\mu}(\cdot|\theta) = \pi(\cdot|\theta)$, i.e., if and only if the elicited likelihoods would coincide with the exogenously given signal. But again, typically we only have access to data of choices among Θ-measurable acts, meaning that we cannot directly elicit the subjective likelihoods.

4. Discussion

4.1. Misspecified beliefs

In recent years, there is literature that aims at understanding when violations of Bayesian updating can be attributed to misspecified beliefs, i.e., to incorrect understanding of the information structure (Shmaya and Yariv, 2016; Bohren and Hauser, 2025, 2026; Molavi, 2026). Notice that the question in these papers bears some analogy to what we do in this paper, i.e., they also try to reconcile elicited beliefs with the information structure. The difference is that they restore consistency by permitting mistakes in the information structure, whereas we restore consistency by permitting mistakes in belief elicitation. In this sense, they attribute potential departures from Bayesian updating to the agent herself, whereas we attribute the same departures to measurement errors due to unaccounted signal dependence of the utilities.

In order to understand how these two approaches relate to each other, let us introduce misspecification in our framework. Accordingly, we first elicit the prior $\bar{\mu}_\Theta$ and the posteriors $\bar{\mu}(\cdot|s)$, and then we ask whether we can find some $\tilde{\pi} : \Theta \to \Delta(S)$ such that

$$\bar{\mu}(\theta|s) = \frac{\bar{\mu}_\Theta(\theta)\tilde{\pi}(s|\theta)}{\sum_{\theta' \in \Theta} \bar{\mu}_\Theta(\theta')\tilde{\pi}(s|\theta')} \tag{5}$$

is satisfied for every $(\theta, s) \in \Omega$.

It is not difficult to see that this is always possible if we set

$$\tilde{\pi}(\cdot|\theta) := \bar{\mu}(\cdot|\theta).$$

In this case, we will obtain $\mathbb{SEU}_{\tilde{\pi}} = \mathbb{SEU}_\Theta$. That is, the only thing that changes is the interpretation of observed violations of Bayesian updating. In our case, we exogenously assume a correct model (i.e., we consider models in $\mathbb{SEU}_\pi$) and we allow for signal-dependent utilities (i.e., we search for a model in $\mathbb{SEU}_\Theta$). On the other hand, in the literature on misspecified beliefs, they implicitly assume signal-independence (i.e., they consider models in $\mathbb{SEU}_\Theta$) and they allow for misspecified models (i.e., they search for a model in $\mathbb{SEU}_{\tilde{\pi}}$).

Of course, the literature on misspecified beliefs focuses on characterizing when Bayesian updating can be rationalized *primarily outside expected utility*, while our paper focuses on rationalizing Bayesian updating *within expected utility*. In this sense, our work can be seen as complementary to the literature on misspecification.

4.2. Nature of beliefs

As we have already discussed, according to the traditional decision-theoretic approach, beliefs do not exist and only acquire meaning within a SEU representation (Grant, 2008; Karni, 2014).

On the other hand, according the alternative view, beliefs are a primitive parameter that describes the agent's uncertainty, i.e., there exists an actual belief (Tsakas, 2026). In this last case, once the actual belief has been pinned down, the actual utility function can be uniquely identified up to strictly positive linear transformations.

Remarkably, for the sake of this paper, this distinction is irrelevant, as our main result distinguishes between two cases, viz., the case where all signal-independent representations are consistent with the information structure (corroborated Bayesian updating), and the case where no signal-independent representation is consistent with the information structure (rejected Bayesian updating). Let us elaborate.

Suppose that we adopt the view that actual beliefs do exist, and use one of the methods from the literature on state-dependent SEU to elicit the actual prior $\hat{\mu}_\Theta$ and the posterior beliefs $\hat{\mu}(\cdot|s)$ for each $s \in S$ (Karni *et al.*, 1983; Lu, 2019; Tsakas, 2026). In this case, it is not difficult to verify that Bayesian updating will be corroborated if and only if the actual SEU representation satisfies $(\hat{u}, \hat{\mu}) \in \mathbb{SEU}_\Theta \cap \mathbb{SEU}_\pi$.

4.3. Dynamic consistency

In the standard decision theoretic literature Bayesian updating is typically understood to always hold if the preferences satisfy the SEU axioms, based on the fact that SEU preferences are dynamically consistent (Karni and Schmeidler, 1991; Epstein and Le Breton, 1993; Ghirardato, 2002).

Formally, dynamic consistency says that if (u, μ) is a SEU representation of $\succeq$, then for any event E the conditional preferences $\succeq_E$ are represented $(u, \mu(\cdot|E))$. That is, conditioning with some event E does not affect the utilities, but rather only the beliefs, which are furthermore updated in a Bayesian way. Obviously, this property is a direct consequence of $\succeq$ satisfying $(A_0) - (A_5)$. So, why is it the case that Bayesian updating can still be rejected?

The reason lies in the fact that in our paper we can only observe choices among Θ-measurable acts, while at the same time we condition with respect to some signal realization $s \in S$, which is not Θ-measurable. As a result, dynamic consistency depends on whether the utility function is Θ-measurable itself.

To understand this last point observe that when evaluation some $f \in \mathcal{F}_\Theta$, we essentially bundle all worlds that share the same θ together to obtain the aggregate state-dependent utility function

$$u_\theta(f_\theta) := \sum_{s \in S} \mu(s|\theta) u_{\theta,s}(f_\theta).$$

This means that the conditional preferences given any Θ-measurable event are dynamically

consistent, i.e., formally, for any $E \subseteq \Theta$ and any two $f, g \in \mathcal{F}_\Theta$, we will have

$$f \succeq_E g \;\Leftrightarrow\; \sum_{\theta\in\Theta} \mu_\Theta(\theta|E)u_\theta(f_\theta) \geq \sum_{\theta\in\Theta} \mu_\Theta(\theta|E)u_\theta(g_\theta).$$

However, conditioning on s, implies that we can no longer use u_θ together with $\mu(\cdot|s)$. Instead we should use $u_{\theta,s}$ together with $\mu(\cdot|s)$. Failure to adjust the utility function will result into rejection of dynamic consistency, and a fortiori of Bayesian updating, even though the preferences over $\mathcal{F}_\Theta$ satisfy the Expected Utility axioms.

A. Belief elicitation mechanisms

The most common elicitation mechanisms include the family of scoring rules (De Finetti, 1937; Brier, 1950; Savage, 1971; Hossain and Okui, 2013; Tsakas, 2020), the family of BDM mechanisms (Karni, 2009; Tsakas, 2019) and the method of matching probabilities (Ducharme and Donnell, 1973; Baillon *et al.*, 2018). For an overview of this literature, see Schlag and van der Weele (2013) and Schotter and Trevino (2014).

As we have already mentioned, all these belief elicitation mechanisms elicit the state-independent belief. Let us illustrate that this is the case for the binarized quadratic scoring rule (Hossain and Okui, 2013).

Accordingly, for a reported marginal belief $r \in \Delta(\Theta)$ and realized state $\theta \in \Theta$, the agent receives a lottery that assigns positive probability to only two outcomes, $x_0, x_1 \in X$ where $x_1 \succ_\omega x_0$ for every $\omega \in \Omega$. In particular, the agent receives x_1 with probability $1 - \|r - e_\theta\|^2$, where e_θ is the Dirac measure that puts probability 1 to θ and $\|\cdot\|$ is the Euclidean distance. Then, take an arbitrary SEU representation (u, μ) and normalize $u_{\theta,s}(x_0) := 0$. This means that the agent will optimally report the marginal belief $r \in \Delta(\Theta)$ that maximizes her SEU,

$$\mathbb{E}_\mu(u(r)) := \sum_{\theta\in\Theta} \mu(\theta)\Big(1 - \|r - e_\theta\|^2\Big) \sum_{s\in S} \mu(s|\theta)u_{\theta,s}(x_1). \tag{A.1}$$

Then, it is not difficult to verify that the unique solution to this maximization problem is to report $r = \bar{\mu}_\Theta$. For the sake of completeness, see Lemma A1 in the Appendix. Of course, the solution does not depend on the choice of the SEU representation as the preferences remain the same.

Lemma A1. *If the actual belief is μ_Θ, the unique optimal report is $\bar{\mu}_\Theta$.*

Proof. The SEU from reporting $r \in \Delta(\Theta)$ can be rewritten as follows:

$$\begin{aligned}
\mathbb{E}_\mu(u(r)) &= \sum_{\theta\in\Theta}\sum_{s\in S}\mu(\theta,s)\Big(1-\|r-e_\theta\|^2\Big)u_{\theta,s}(x_1)\\
&= \sum_{\theta\in\Theta}\sum_{s\in S}\bar{\mu}(\theta,s)\Big(1-\|r-e_\theta\|^2\Big)\bar{u}(x_1)\\
&= \bar{u}(x_1)\sum_{\theta\in\Theta}\bar{\mu}(\theta)\Big(1-\|r-e_\theta\|^2\Big).
\end{aligned}$$

Hence, it suffices to show

$$\begin{aligned}
\mathbb{E}_\mu(u(\bar{\mu}_\Theta)) > \mathbb{E}_\mu(u(r)) &\Leftrightarrow \sum_{\theta\in\Theta}\bar{\mu}_\Theta(\theta)\Big(1-\|\bar{\mu}_\Theta-e_\theta\|^2\Big) > \sum_{\theta\in\Theta}\bar{\mu}_\Theta(\theta)\Big(1-\|r-e_\theta\|^2\Big)\\
&\Leftrightarrow \sum_{\theta\in\Theta}\bar{\mu}_\Theta(\theta)\Big(2\langle\bar{\mu}_\Theta,e_\theta\rangle-\|\bar{\mu}_\Theta\|^2\Big) > \sum_{\theta\in\Theta}\bar{\mu}_\Theta(\theta)\Big(2\langle r,e_\theta\rangle-\|r\|^2\Big)\\
&\Leftrightarrow \sum_{\theta\in\Theta}\bar{\mu}_\Theta(\theta)\Big(2\bar{\mu}_\Theta(\theta)-\|\bar{\mu}_\Theta\|^2\Big) > \sum_{\theta\in\Theta}\bar{\mu}_\Theta(\theta)\Big(2r(\theta)-\|r\|^2\Big)\\
&\Leftrightarrow \|\bar{\mu}_\Theta\|^2 > 2\langle\bar{\mu}_\Theta,r\rangle-\|r\|^2\\
&\Leftrightarrow \|\bar{\mu}_\Theta-r\| > 0.
\end{aligned}$$

Obviously the strict inequality holds for every $r \neq \bar{\mu}_\Theta$, which completes the proof. □

B. Proofs

Lemma B2. *For a full-support $\mu \in \Delta(\Omega)$, if*

$$\frac{\mu(\theta_1|s)}{\mu(\theta_2|s)} = \frac{\mu(\theta_1)}{\mu(\theta_2)}\frac{\pi(s|\theta_1)}{\pi(s|\theta_2)} \tag{B.1}$$

holds for every $\theta_1,\theta_2 \in \Theta$ and $s \in S$, then it is the case that $\mu \in \mathcal{M}_\pi$.

Proof. It is not difficult to verify that (B.1) implies that for every $s \in S$ there is some $d_s > 0$ such that $\mu(s|\theta) = d_s\pi(s|\theta)$ for all (θ,s). Thus, we obtain

$$\sum_{s\in S} d_s\pi(s|\theta) = 1$$

for all $\theta \in \Theta$. In matrix language this means $\Pi d = 1_\Theta$. By the fact that Π is full-rank, there is a left inverse Q. In addition, by Π being row-stochastic, the entries of each row of Q add up to 1. Hence, we obtain

$$d = Q1_\Theta = 1_S,$$

which directly implies $\mu(s|\theta) = \pi(s|\theta)$ for all (θ,s), meaning that $\mu \in \mathcal{M}_\pi$. □

Lemma B3. *Suppose that* $\succeq$ *satisfies* $(A_0)-(A_5)$. *Then, we have either* $\mathbb{SEU}_\Theta \cap \mathbb{SEU}_\pi = \emptyset$ *or* $\mathbb{SEU}_\Theta = \mathbb{SEU}_\pi$.

Proof. Suppose that there exists some $(u,\mu) \in \mathbb{SEU}_\Theta \cap \mathbb{SEU}_\pi$. By $(u,\mu) \in \mathbb{SEU}_\Theta$, it is the case that the multiplier parameter $\beta_{\theta,s}$ is Θ-measurable, i.e., for any (θ, s),

$$\frac{\bar{\mu}(\theta,s)}{\mu(\theta,s)} = \beta_\theta.$$

By $(u,\mu) \in \mathbb{SEU}_\pi$, we obtain

$$\frac{\mu(\theta_1|s)}{\mu(\theta_2|s)} = \frac{\mu(\theta_1)}{\mu(\theta_2)}\frac{\pi(s|\theta_1)}{\pi(s|\theta_2)}$$

for every $\theta_1, \theta_2 \in \Theta$ and $s \in S$. By multiplying both sides with $\beta_{\theta_2}/\beta_{\theta_1}$, we obtain

$$\frac{\bar{\mu}(\theta_1|s)}{\bar{\mu}(\theta_2|s)} = \frac{\bar{\mu}(\theta_1)}{\bar{\mu}(\theta_2)}\frac{\pi(s|\theta_1)}{\pi(s|\theta_2)}, \tag{B.2}$$

meaning that $(\bar{u},\bar{\mu}) \in \mathbb{SEU}_\pi$. Then, we proceed to prove the following two claims:

Claim 1. $\mathbb{SEU}_\Theta \subseteq \mathbb{SEU}_\pi$.

Take some $(\tilde{u},\tilde{\mu}) \in \mathbb{SEU}_\Theta$, meaning that the corresponding multiplier parameter $\tilde{\beta}_{\theta,s}$ is also Θ-measurable. Then, by multiplying both sides of (B.2) with $\tilde{\beta}_{\theta_1}/\tilde{\beta}_{\theta_2}$, we obtain

$$\frac{\tilde{\mu}(\theta_1|s)}{\tilde{\mu}(\theta_2|s)} = \frac{\tilde{\mu}(\theta_1)}{\tilde{\mu}(\theta_2)}\frac{\pi(s|\theta_1)}{\pi(s|\theta_2)},$$

meaning that $(\tilde{u},\tilde{\mu}) \in \mathbb{SEU}_\pi$.

Claim 2. $\mathbb{SEU}_\pi \subseteq \mathbb{SEU}_\Theta$.

Take some $(\tilde{u},\tilde{\mu}) \in \mathbb{SEU}_\pi$. Then, we obtain

$$\frac{\bar{\mu}(\theta_1,s)}{\bar{\mu}(\theta_2,s)}\frac{\bar{\mu}(\theta_1)}{\bar{\mu}(\theta_2)} = \frac{\tilde{\mu}(\theta_1,s)}{\tilde{\mu}(\theta_2,s)}\frac{\tilde{\mu}(\theta_1)}{\tilde{\mu}(\theta_2)} = \frac{\tilde{\beta}_{\theta_2,s}}{\tilde{\beta}_{\theta_1,s}}\frac{\bar{\mu}(\theta_1,s)}{\bar{\mu}(\theta_2,s)}\frac{\tilde{\mu}(\theta_1)}{\tilde{\mu}(\theta_2)},$$

where the first equality follows from (B.2) combined with the fact that $(\tilde{u},\tilde{\mu}) \in \mathbb{SEU}_\pi$, and the second equality follows by construction of the multiplier parameter. This directly implies

$$\frac{\tilde{\beta}_{\theta_1,s}}{\tilde{\beta}_{\theta_2,s}} = \frac{\bar{\mu}(\theta_1)}{\bar{\mu}(\theta_2)}\frac{\tilde{\mu}(\theta_2)}{\tilde{\mu}(\theta_1)}. \tag{B.3}$$

The latter implies that the $|\Theta| \times |S|$ matrix $\tilde{B}$ with entries $\tilde{\beta}_\theta$ has rank equal to 1. Therefore there are two vectors $(b_\theta)_{\theta\in\Theta}$ and $(c_s)_{s\in S}$ with strictly positive entries, such that for every $(\theta,s) \in \Omega$ we obtain

$$\tilde{\beta}_{\theta,s} = b_\theta c_s. \tag{B.4}$$

Combining the latter with Equation (B.3) yields

$$\begin{aligned}\frac{b_{\theta_1}}{b_{\theta_2}} &= \frac{\bar{\mu}(\theta_1)}{\bar{\mu}(\theta_2)}\frac{\tilde{\mu}(\theta_2)}{\tilde{\mu}(\theta_1)} \\ &= \frac{\sum_{s\in S}\bar{\mu}(\theta_1,s)}{\sum_{s\in S}\bar{\mu}(\theta_2,s)}\frac{\tilde{\mu}(\theta_2)}{\tilde{\mu}(\theta_1)} \\ &= \frac{b_{\theta_1}}{b_{\theta_2}}\frac{\sum_{s\in S}c_s\tilde{\mu}(\theta_1,s)}{\sum_{s\in S}c_s\tilde{\mu}(\theta_2,s)}\frac{\tilde{\mu}(\theta_2)}{\tilde{\mu}(\theta_1)} \\ &= \frac{b_{\theta_1}}{b_{\theta_2}}\frac{\sum_{s\in S}c_s\tilde{\mu}(s|\theta_1)}{\sum_{s\in S}c_s\tilde{\mu}(s|\theta_2)},\end{aligned}$$

meaning that for all $\theta \in \Theta$,

$$\sum_{s\in S} c_s\pi(s|\theta) = \gamma > 0. \tag{B.5}$$

Hence, using matrix language, we obtain

$$\Pi c = \gamma 1_\Theta.$$

Once again, since Π is a row stochastic matrix with full rank, we can take the left inverse Q whose every row adds up to 1, which we already used in the proof of Lemma B2. As a result, we obtain $c = \gamma Q 1_\Theta$, meaning that

$$c_s = \gamma \tag{B.6}$$

for every $s \in S$. Therefore, it is the case that $\tilde{\beta}_{\theta,s} = b_\theta$ for all $s \in S$ and all $\theta \in \Theta$, which in turn implies that $(\tilde{u}, \tilde{\mu}) \in \mathbb{SEU}_\Theta$.

Finally, putting the two claims together yields $\mathbb{SEU}_\Theta = \mathbb{SEU}_\pi$, which completes the proof. □

Proof of Theorem 1. By Lemma B2, Bayesian updating is correborated if and only if $(\bar{u}, \bar{\mu}) \in \mathbb{SEU}_\pi$. Then, by Lemma B3, we have that $(\bar{u}, \bar{\mu}) \in \mathbb{SEU}_\pi$ if and only if $\mathbb{SEU}_\pi = \mathbb{SEU}_\Theta$, while $(\bar{u}, \bar{\mu}) \notin \mathbb{SEU}_\pi$ if and only if $\mathbb{SEU}_\pi \cap \mathbb{SEU}_\Theta = \emptyset$, which completes the proof. □

References

Anscombe, F. and Aumann, R. (1963). A definition of subjective probability. *Annals of Mathematical Statistics*, **34**, 199–205.

Baillon, A., Huang, Z., Selim, A. and Wakker, P. (2018). Measuring ambiguity attitudes for all (natural) events. *Econometrica*, **86**, 1839–1858.

Benjamin, D. (2019). Errors in probabilistic reasoning and judgment biases. *Handbook of Behavioral Economics – Foundations and Applications*, **2**, 69–186.

Bohren, A. and Hauser, D. (2025). Misspecified models in learning and games. *Annual Review of Economics*, **17**.

— and — (2026). Behavioral foundations of model misspecification. *Working Paper*.

Brier, G. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, **78**, 1–3.

Caplin, A. and Leahy, J. (2001). Psychological expected utility theory and anticipatory feelings. *Quarterly Journal of Economics*, **116**, 55–79.

De Finetti, B. (1937). La prévision: ses lois logiques, ses sources subjectives. *Annales de Ínstitut Henri Poincaré*, **7**, 1–78.

Ducharme, W. and Donnell, M. (1973). Intrasubject comparison of four response modes for subjective probability assessment. *Organizational Behavior and Human Performance*, **10**, 108–117.

Epstein, L. and Le Breton, M. (1993). Dynamically consistent beliefs must be bayesian. *Journal of Economic Theory*, **61**, 1–22.

Ghirardato, P. (2002). Revisiting savage in a conditional world. *Economic Theory*, **20**, 83–92.

Golman, R., Loewenstein, G., Molnar, A. and Saccardo, S. (2022). The demand for, and avoidance of, information. *Management Science*, **68**, 6355–7064.

Grant, T., Sion & van Zandt (2008). Expected utility theory. *Handbook of Rational and Social Choice*, pp. 21–68.

Hossain, T. and Okui, R. (2013). The binarized scoring rule. *Review of Economic Studies*, **80**, 984–1001.

Karni, E. (2009). A mechanism for eliciting probabilities. *Econometrica*, **77**, 603–606.

— (2014). Axiomatic foundations of expected utility and subjective probability. *Handbook of the Economics of Risk and Uncertainty*, **1**, 1–39.

— and Schmeidler, D. (1991). Atemporal dynamic consistency and expected utility theory. *Journal of Economic Theory*, **54**, 401–408.

—, — and Vind, K. (1983). On state-dependent preferences and subjective probabilities. *Econometrica*, **51**, 1021–1031.

Köszegi, B. (2006). Emotional agency. *Quarterly Journal of Economics*, **121**, 121–155.

Köszegi, B. and Rabin, M. (2009). Reference-dependent consumption plans. *American Economic Review*, **99**, 909–936.

Lu, J. (2019). Bayesian identification: a theory for state-dependent utilities. *American Economic Review*, **109**, 3192–3228.

Molavi, P. (2026). The empirical content of bayesian updating under misspecification. *Working Paper*.

Savage, L. (1971). Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, **66**, 783–801.

Schlag, K. and van der Weele, J. (2013). Eliciting probabilities, means, medians, variances and covariances without assuming risk neutrality. *Theoretical Economics Letters*, **3**, 38–42.

Schotter, A. and Trevino, I. (2014). Belief elicitation in the laboratory. *Annual Review of Economics*, **6**, 103–128.

Shmaya, E. and Yariv, L. (2016). Experiments on decisions under uncertainty: A theoretical framework. *American Economic Review*, **106**, 1775–1801.

Tsakas, E. (2019). Obvious belief elicitation. *Games and Economic Behavior*, **118**, 374–381.

— (2020). Robust scoring rules. *Theoretical Economics*, **15**, 955–987.

— (2026). Belief identification by proxy. *Review of Economic Studies*, **93**, 697–724.