University of Pretoria
Department of Economics Working Paper Series

Bayesian Learning with Multiple Priors and Non-Vanishing Ambiguity

Alexander Zimper
University of Pretoria
Wei Ma
University of Pretoria
Working Paper: 2015-35
June 2015

Department of Economics
University of Pretoria
0002, Pretoria
South Africa
Tel: +27 12 420 2413

Bayesian learning with multiple priors and non-vanishing ambiguity*

Alexander Zimper† Wei Ma‡

June 1, 2015

Abstract

The existing models of Bayesian learning with multiple priors by Marinacci (2002) and by Epstein and Schneider (2007) formalize the intuitive notion that ambiguity should vanish through statistical learning in an one-urn environment. Moreover, the multiple priors decision maker of these models will eventually learn the 'truth'. To accommodate non vanishing violations of Savage's (1954) sure-thing principle, as reported in Nicholls et al. (2015), we construct and analyze a model of Bayesian learning with multiple priors for which ambiguity does not necessarily vanish. Our decision maker only forms posteriors from priors that pass a plausibility test in the light of the observed data in the form of a γ-maximum expected loglikelihood prior-selection rule. The "stubbornness" parameter $\gamma \geq 1$ determines the magnitude by which the expectation of the loglikelihood with respect to plausible priors can differ from the maximal expected loglikelihood. The greater the value of γ, the more priors pass the plausibility test to the effect that less ambiguity vanishes in the limit of our learning model.

JEL Classification: C11, D81.

Keywords: Ambiguity, Bayesian Learning, Misspecified Priors, Berk's Theorem, Kullback-Leibler Divergence, Ellsberg Paradox

---

*We thank Larry Epstein, Massimo Marinacci, and an anonymous referee for helpful comments and suggestions. Financial support from ERSA (Economic Research Southern Africa) is gratefully acknowledged.

†Corresponding author. Department of Economics, University of Pretoria, Private Bag X20, Hatfield 0028, South Africa. Tel: 0027 12 420-3522. Fax: 0027 86 691-2749. E-mail: alexander.zimper@up.ac.za

‡Department of Economics, University of Pretoria, South Africa. E-mail: maweicityu@gmail.com

1 Introduction

In a seminal contribution, Savage (1954) provides an axiomatic foundation for *subjective expected utility* (SEU) theory which resolves a decision maker's uncertainty through a unique additive (subjective) probability measure. However, starting with Ellsberg's (1961) one-urn experiment, several experimental studies report systematic violations of Savage's key axiom, the *sure-thing principle* (STP) (cf. Wu and Gonzales 1999; Wakker 2010 and references therein). As a reaction to this finding, descriptive decision theories have been developed which explain violations of the STP through ambiguity attitudes. Central to this paper are multiple priors models which use sets of subjective additive probability measures rather than a unique measure to describe a decision maker's uncertainty (cf. Gilboa and Schmeidler 1989; Jaffray 1994; Ghirardato, Maccheroni, and Marinacci 2004).[1] Multiple priors models offer a straightforward interpretation of ambiguity as a lack of 'probabilistic' information: "[...] the subject has too little information to form a prior. Hence (s)he considers a set of priors." (Gilboa and Schmeidler 1989, p. 142). By this interpretation, one would intuitively expect that violations of the STP must vanish if the decision maker observes an unlimited amount of statistical information. Our intuition is thereby informed by standard models of Bayesian learning according to which a Savage (1954) decision maker–who holds a unique subjective prior–will (under some regularity condition) almost certainly learn the 'true' probability measure if he observes a large amount of data which was i.i.d. generated by this measure.

A recent experimental study by Nicholls, Romm, and Zimper (2015) has put the notion to the test that STP violations should tend to decrease through statistical learning. These authors were running a sequence of Ellsberg-type one-urn experiments such that the test group received an increasing amount of statistical information about the urn's true composition whereas the control group did not receive such information. Quite surprisingly, the authors find that "... statistical learning has, at best, no impact on STP violations. At worst, it might even be causing STP violations to increase." (Nicholls et al. 2015, p. 14)

To accommodate this empirical finding, this paper constructs a model of Bayesian learning with multiple priors such that ambiguity does not necessarily vanish through statistical learning. As the key feature of our model, the decision maker rejects priors in the light of observed data by an application of the *γ-maximum expected loglikelihood prior-selection rule*, which we newly introduce to the literature. The remainder of this

[1] An alternative (and under specific circumstances formally equivalent) class of models that accommodate ambiguity attitudes are models of Choquet decision making/Choquet expected utility (Schmeidler 1989; Gilboa 1987). These Choquet models express ambiguity attitudes through non-additive probability measures.

introduction explains our formal learning model, as well as its economic relevance, in more detail.

1.1 Existing results on Bayesian learning

Consider an indexed family of probability measures such that, unbeknownst to the decision maker, one measure in this family (e.g., given by the composition of an urn) is the true data-generating measure. The standard model of Bayesian learning considers a Savage (1954) decision maker who resolves his uncertainty about the true measure through a unique prior defined on an index space. In the one-urn environment relevant to this paper, the index space is in an one-one relationship with the family of measures. When this decision maker observes an i.i.d. data sample generated by the true measure, he uses this statistical information to update his prior to a posterior by an application of Bayes' rule. If the prior is *well-specified*, i.e., if the true index belongs to its support, standard consistency results imply that the decision maker's posteriors will almost surely converge towards a Dirac measure concentrating at the true index/measure when he can observe an unlimited amount of statistical information.[2] More generally, for *well-* and *misspecified* priors, Berk's (1966) theorem implies that the posteriors will almost surely concentrate at the index in the prior's support that minimizes the *Kullback-Leibler (1951) divergence*[3] from the true measure.

Turn now to a multiple priors decision maker who resolves his uncertainty about the true index/measure by a set of priors rather than a unique prior. Existing formal models of Bayesian learning with multiple priors by Marinacci (2002) (=M-2002) and by Epstein and Schneider (2007) (=ES-2007) establish formal conditions such that all multiple posteriors concentrate at the true index/measure. Under the assumptions of these models, STP-violations will thus vanish through Bayesian learning in the single likelihood environment relevant to the Ellsberg one-urn experiment. More specifically, M-2002 proves convergence to the true index/measure under the assumption that all priors are well-specified. ES-2007 assume that the decision maker applies a specific prior-selection rule–which we call the α*-expected maximum likelihood* rule–according to which he rejects priors that are implausible in the light of the observed data. Posteriors are then only formed from priors that are not rejected. Restricted to the one-urn environment, ES-2007's Theorem 1 implies that all multiple posteriors will concentrate at the true index/measure if there is (at least) one well-specified prior.

[2]The seminal contribution is Doob's (1949) consistency theorem. For generalizations and further references see, e.g. Diaconis and Freedman (1986), Chapter 1 in Gosh and Ramamoorthi (2003), Lijoi, Pruenster and Walker (2004).

[3]For a formal definition see Section 2.2.

1.2 The γ-maximum expected loglikelihood rule

In the M-2002 model, the decision maker forms posteriors from all priors. In contrast, ES-2007 assume that the decision maker only forms posteriors from priors that pass a plausibility test in the form of their α-expected maximum likelihood rule. On the one hand, we follow ES-2007 in that we regard it as plausible that a multiple priors decision maker should reject priors that are, by some plausibility criterion, at odds with the observed data. On the other hand, the α-expected maximum likelihood rule might be too strong for some multiple priors decision makers because it implies vanishing ambiguity in the single-urn environment.[4] Because we want to establish the possibility of non-vanishing STP violations, we introduce the *γ-maximum expected loglikelihood* rule as a plausible alternative to the α-expected maximum likelihood rule.

Formally, ES-2007's α-expected maximum likelihood rule rejects priors whose expected likelihood for a given data sample is not α-close to the maximal expected likelihood for some fixed parameter $\alpha \in (0, 1]$. In contrast, our γ-maximum expected loglikelihood rule rejects, for a fixed $\gamma \in [1, \infty)$, priors that are not γ-close to the maximal expected *log*likelihood. Both rules are equivalent if all priors are degenerate (i.e., Dirac) probability measures since likelihood and loglikelihood maximization are identical. However, if expectations of likelihoods versus loglikelihoods are taken with respect to non-degenerate priors, the γ-maximum expected loglikelihood rule punishes more strongly priors that support indices with small likelihoods. We therefore interpret our decision maker as more cautious (i.e., more risk averse with respect to likelihood outcomes) than the ES-2007 decision maker.

Two main findings for Bayesian learning with multiple priors under the γ-maximum expected loglikelihood rule emerge.

1. For the special case of the maximum expected loglikelihood rule (i.e., $\gamma = 1$), all multiple posteriors concentrate at the (typically) unique index/measure that minimizes the Kullback-Leibler divergence from the true measure over all indices in the support of the prior that minimizes the *expected* Kullback-Leibler divergence from the true measure. In contrast to ES-2007's α-expected maximum likelihood rule, however, this unique index/measure is not necessarily the true index/measure even if there is a well-specified prior.

[4]There might exist a number of possible 'explanations' for the experimental findings of Nicholls et al. (2015). However, under the assumption that the subjects resemble Bayesian learners with multiple priors, we would–contrary to the experimental findings–expect STP violations to decrease under the ES-2007 learning model. In contrast, a learning model with non-vanishing ambiguity could more convincingly explain persistent STP violations for this single-urn/likelihood experiment.

2. Larger values of $\gamma > 1$ lead to larger sets of posteriors whereby we can always find sufficiently large values of γ such that (for sufficiently rich sets of priors) some posterior will concentrate at any given index/measure.

1.3 Non-vanishing violations of the sure-thing principle

To illustrate the possibility of non-vanishing STP violations, we reformulate the original Ellsberg (1961) one-urn experiment within our framework of Bayesian learning with multiple priors. We distinguish between an *a priori* (i.e., before any statistical information has been observed) and an *a posteriori* (i.e., after an unlimited amount of statistical information has been observed) one-shot decision situation. We speak of non-vanishing STP violations if the decision maker commits an Ellsberg paradox in the *a priori* as well as in the *a posteriori* decision situation. Under the assumption that the multiple priors decision maker is a *maxmin expected utility* decision maker (cf. Gilboa and Schmeidler 1989), we identify conditions such that our learning model gives rise to non-vanishing STP violations. More concretely, we show how the set of posteriors that emerge in the limit of the Bayesian learning process gradually increases (with respect to set-inclusion) in the value of the γ parameter. In other words, non-vanishing ambiguity is γ-sensitive. This parameter sensibility of non-vanishing ambiguity is in contrast to ES-2007's α-expected maximum likelihood rule where ambiguity completely vanishes for any parameter value $\alpha > 0$.

1.4 Economic relevance

The possibility of non-vanishing ambiguity is a potentially attractive feature in economic applications. Consider, for example, the class of theoretical models that establish the possibility of speculative trade under the assumption that the decision makers express ambiguity attitudes (e.g., Dow, Madrigal, and Werlang 1990; Halevy 2004; Zimper 2009; Werner 2014). In contrast to the speculative trade model of Harrison and Kreps (1978), which is based on heterogenous additive beliefs, speculative trade in these ambiguity-driven models might become persistent under non-vanishing ambiguity even if the agents are Bayesian learners.

As another example, consider macro-economic models which deviate from Muth's (1961) rational expectations paradigm. Here, our model's parameter-sensibility of non-vanishing ambiguity is especially relevant. Similar to different values of a personal risk-aversion parameter (as, e.g., in CRRA or CARA utility functions), different values of the γ-parameter can be used to describe a personal feature of economic agents.[5] In dynamic

[5]Daniele Pennesi suggested to call the parameter γ a "stubbornness measure" because it reflects the

models where multiple priors agents update their priors in the light of a large amount of statistical information, agents with small values of γ will closely resemble a rational expectations EU decision maker whereas agents with large values of γ might express strong ambiguity attitudes. The γ-sensibility of non-vanishing ambiguity thus admits for a comparative statics analysis or/and for heterogenous agents models. For example, the 'risk-free rate' and the 'equity premium' puzzles put forward by Mehra and Prescott (1985; 2003) are based on the assumption that the representative agent's belief about the consumption growth rate coincides with its objective distribution. This assumption is in turn justified by existing consistency results for Bayesian learning with single as well as with multiple prior(s) combined with the large amount of statistical data on consumption growth available to the representative agent. We consider it an interesting avenue for future research to investigate in how far multiple priors decision making embedded into our Bayesian learning model might contribute towards an explanation of these asset pricing puzzles for plausible values of γ.

The remainder of the paper is organized as follows. Section 2 formally introduces the one-urn environment and recalls Berk's (1966) Theorem. Section 3 extends Bayesian learning to the multiple priors framework. In Section 4 we present our main formal results for Bayesian learning with multiple priors under the γ-maximum expected loglikelihood prior-selection rule. Section 5 applies our theoretical findings to non-vanishing STP violations in Ellsberg's (1961) one-urn experiment. Section 6 discusses possible extensions of our approach in comparison with the existing literature on Bayesian learning under ambiguity. Section 7 concludes. All formal proofs are relegated to the Appendix.

2 Preliminaries

2.1 Set-up: The one-urn environment

Denote by (Ω, Σ) a measurable space with state space Ω and σ-algebra Σ. This paper is exclusively concerned with two special cases of measurable spaces. First, we speak of the *continuous case* if Ω is some subset of the Euclidean line $\mathbb{R}$ and Σ is the corresponding Borel σ-algebra. Second, we speak of the *finite case* whenever Ω is finite with $\#\Omega > 1$ and Σ is the power-set of Ω.

Unbeknownst to the decision maker, there exists a 'true'/'objective' probability measure defined on (Ω, Σ), denoted φ_{θ^*}. To capture this lack of knowledge, we consider a

decision maker's lack of willingness to revisit his priors in the light of new data.

set of probability measures on (Ω, Σ) that are indexed by $\theta \in \Theta$, i.e.,

$$\Phi = \{\varphi_\theta \mid \theta \in \Theta\}. \tag{1}$$

We assume that Θ is finite with $\#\Theta = n \geq 2$ and that $\theta^* \in \Theta$. Because we want to avoid Bayesian updating in the light of events that were ex ante perceived as impossible, we further assume that all measures in Φ have full support on Ω. Denote by $(\Theta, \mathcal{F})$ the index space such that $\mathcal{F}$ is the powerset of Θ.[6]

Next we consider a Σ-measurable function (random variable) $X : \Omega \to \mathbb{R}$ such that, for all $\varphi_\theta \in \Phi$ and all $B \in \Sigma$,

$$\begin{aligned} E\left[I_{X^{-1}(B)}, \varphi_\theta\right] &= \int_\Omega I_B(\omega)\, d\varphi_\theta(\omega) && (2) \\ &= \varphi_\theta(B) && (3) \end{aligned}$$

where I_B denotes the indicator function of B. Since the distribution function (=cdf) of X on the probability space $(\Omega, \Sigma, \varphi_\theta)$ fully specifies the measure φ_θ, we slightly abuse notation by identifying X's cdf, denoted $\varphi_\theta : \mathbb{R} \to [0,1]$, on $(\Omega, \Sigma, \varphi_\theta)$ with the corresponding probability measure $\varphi_\theta : \Sigma \to [0,1]$. That is, X satisfies, for any $\varphi_\theta \in \Phi$,

$$\varphi_\theta((a,b]) = \varphi_\theta(X = b) - \varphi_\theta(X = a) \text{ for all } (a,b] \in \Sigma, \tag{4}$$

in the continuous case; and

$$\varphi_\theta(\{\omega_m\}) = \varphi_\theta(X = b) - \varphi_\theta(X = a) \text{ for all } \omega_m \in \Omega \tag{5}$$

with $X(\omega_m) \leq b < X(\omega_{m+1})$, $X(\omega_{m-1}) \leq a < X(\omega_m)$ in the finite case, respectively.

By the above set-up, the decision maker's uncertainty about the true probability measure in Φ on (Ω, Σ) is equivalent to his uncertainty about the true distribution of X. Furthermore, both notions of uncertainty are equivalent to the decision maker's uncertainty about the true index in Θ. We refer to this one-one correspondence between probability measures and indices as the "one-urn" or "single-likelihood" environment.[7]

We will frequently use the *Radon-Nikodym derivative* of measure φ_θ with respect to a dominating measure m on (Ω, Σ), denoted $\frac{d\varphi_\theta}{dm}$. This derivative is defined such that, for all $B \in \Sigma$,

$$\varphi_\theta(B) = \int_{\omega \in B} \frac{d\varphi_\theta}{dm}(\omega)\, dm. \tag{6}$$

[6]In the literature, $(\Theta, \mathcal{F})$ is also called the (possibly multiple) parameter space.

[7]As a generalization of the single-likelihood environment, ES-2007 consider a "multiple-likelihoods" environment where an index θ in Θ corresponds to a set of θ-conditional probability measures. Although the formal results of this paper will be exclusively derived for the single-likelihood environment, compare Section 6 for an outlook on future research.

In the continuous case, we assume that m is given as the Lebesgue measure so that, for any absolutely continuous distribution function φ_θ, $\frac{d\varphi_\theta}{dm} : \mathbb{R} \rightarrow \mathbb{R}_+$ stands for the familiar probability density function (=pdf) such that

$$\varphi_\theta\left((a,b]\right) = \int_{x\in(a,b]} \frac{d\varphi_\theta}{dm}(x)\, dx. \tag{7}$$

In the finite case, we assume that m is given as the counting measure, implying, for all $\omega \in \Omega$,

$$\begin{aligned} \varphi_\theta(\omega) &= \int_{\omega'\in\{\omega\}} \frac{d\varphi_\theta}{1}(\omega')\, 1 & (8)\\ &= \frac{d\varphi_\theta}{1}(\omega) & (9)\\ &= d\varphi_\theta(\omega). & (10)\end{aligned}$$

In the finite case, $\frac{d\varphi_\theta}{dm}(\omega)$ as well as $d\varphi_\theta(\omega)$ thus become equivalent notions for the probability $\varphi_\theta(\{\omega\})$ of the singleton event $\{\omega\}$.

Example 1. Continuous case "Family of normal distributions". Let $\Omega = \mathbb{R}$ and $X(\omega) = \omega$. Suppose that the probability measures φ_θ in Φ are specified by the cdf's of a normal distribution $\mathcal{N}(\mu_\theta, \sigma_\theta)$ with mean μ_θ and standard deviation σ_θ, $\theta \in \Theta$. That is, the decision maker's uncertainty about the true measure φ_{θ^*} in Φ is equivalent to his uncertainty about the true normal distribution $\mathcal{N}(\mu_{\theta^*}, \sigma_{\theta^*})$ which, in turn, is equivalent to his uncertainty about the true index $\theta^* \in \Theta$. The Radon-Nikodym derivative $\frac{d\varphi_\theta}{dm}$ is here the pdf of $\mathcal{N}(\mu_\theta, \sigma_\theta)$.□

Example 2. Finite case "Coin tossing". Let $\Omega = \{\omega_0, \omega_1\}$ and $X(\omega_k) = k$. Further suppose that

$$\begin{aligned} \omega_0 &= Heads\\ \omega_1 &= Tails \end{aligned}$$

and

$$\varphi_\theta(X=1) = \varphi_\theta(\omega_1) = \theta. \tag{11}$$

Here the decision maker's uncertainty about the true probability measure φ_{θ^*} in Φ is equivalent to his uncertainty about the true probability θ^* of the event $\{Tails\}$ resulting from a coin toss. The index set Θ thus contains the parameters of a Bernoulli distribution. The Radon-Nikodym derivative $\frac{d\varphi_\theta}{dm}(\omega)$, as well as $d\varphi_\theta(\omega)$, gives the probability of event $\{\omega\}$, $\omega \in \Omega$.□

In a next step, we assume that the decision maker can observe data generated by a sequence of independently φ_{θ^*}-distributed coordinate random variables $X_1, X_2, ...$ defined on the probability space $(\Omega^\infty, \Sigma^\infty, P_{\theta^*})$ such that $\Omega^\infty = \times_{t=1}^\infty \Omega$; Σ^∞ denotes the standard product algebra generated by $\Sigma, \Sigma, ...$; and P_{θ^*} is the product measure generated by the φ_{θ^*}'s. Each $X_t : \Omega^\infty \to \mathbb{R}$ is thereby a time t version of the Σ-measurable function $X : \Omega \to \mathbb{R}$ in the sense that

$$X_t\left(..., \omega_t, ...\right) = X\left(\omega\right) \text{ for } \omega_t = \omega. \tag{12}$$

$(\Omega^\infty, \Sigma^\infty)$ is called the sample space because every realization of $X_1, X_2, ...$ corresponds to a data sample that might be possibly observed by the decision maker.

In the absence of ambiguity, the decision maker's uncertainty about the true probability measure on (Ω, Σ) is modeled through a unique additive probability measure–"the prior"–defined on the index space $(\Theta, \mathcal{F})$. In contrast, ambiguity with respect to the true probability measure on (Ω, Σ) will be modeled through a non-degenerate set of additive probability measures–"multiple priors"–defined on the index space $(\Theta, \mathcal{F})$. Models of Bayesian learning investigate how the decision maker forms posteriors from his prior(s) in the light of new statistical information drawn from the sample space.

2.2 Bayesian learning with a unique prior

Consider a standard Bayesian decision maker who holds a unique prior $\mu_0 \in \triangle^n$ defined on the parameter space $(\Theta, \mathcal{F})$ where $\triangle^n$ denotes the Euclidean n-simplex.[8] Through Bayesian updating we obtain the (conditional) probability space $\left(\Theta, \mathcal{F}, \pi_{\mu_0}^t\right)$ such that one version of the posterior $\pi_{\mu_0}^t \in \triangle^n$, formed from the prior μ_0 after observing a data sample drawn from $X_1, ..., X_t$, is formally given as

$$\pi_{\mu_0}^t\left(\Theta'\right) = \frac{\sum_{\theta\in\Theta'} \prod_{i=1}^{t} \frac{d\varphi_\theta}{dm}\left(X_i\right) \cdot \mu_0\left(\theta\right)}{\sum_{\theta\in\Theta} \prod_{i=1}^{t} \frac{d\varphi_\theta}{dm}\left(X_i\right) \cdot \mu_0\left(\theta\right)} \tag{13}$$

[8] Since there is an one-one correspondence between all probability measures on $(\Theta, \mathcal{F})$ and the points in $\triangle^n$, we slightly abuse notation and write $\mu_0 \equiv \left(\mu_0^1, ..., \mu_0^n\right) \in \triangle^n$ for the additive probability measure $\mu_0 : \mathcal{F} \to [0, 1]$ such that, for all non-empty $\Theta' \in \mathcal{F}$,

$$\mu_0\left(\Theta'\right) = \Sigma_{\{\theta_j\in\Theta'\}}\mu_0^j.$$

for any $\Theta' \in \mathcal{F}$. Recall that, in the continuous case, $\frac{d\varphi_\theta}{dm}(X_i = x)$ denotes the evaluated pdf $\frac{d\varphi_\theta}{dm}(x)$ whereas, in the finite case, $\frac{d\varphi_\theta}{dm}(X_i(\omega) = x)$ denotes the probability of state ω with respect to measure φ_θ. Note that, by the martingale convergence theorem, the posterior $\pi^t_{\mu_0}$ converges with probability one to some *emerging posterior* $\pi^\infty_{\mu_0}$.[9]

A prior μ_0 is well-specified if, and only if, the true parameter belongs to the support of μ_0, i.e., for our finite Θ, iff $\mu_0(\theta^*) > 0$. Denote by $\delta_\theta \in \triangle^n$ the Dirac measure that attaches probability one to the index value $\theta \in \Theta$. By Doob's (1949) consistency theorem[10], the emerging posterior of a well-specified prior will almost surely concentrate at the true parameter value if the number t of observations becomes arbitrarily large, i.e.,

$$\pi^\infty_{\mu_0} = \delta_{\theta^*}, \text{ a.s. } P_{\theta^*} \tag{14}$$

or, equivalently,

$$Support\left(\pi^\infty_{\mu_0}\right) = \{\theta^*\}, \text{ a.s. } P_{\theta^*}. \tag{15}$$

To state a convergence result for the more general case of not necessarily well-specified priors, let us recall the following definition due to Kullback and Leibler (1951).

Definition 1. The *Kullback-Leibler (KL) divergence* of φ' from φ is defined as

$$\begin{aligned} D_{KL}(\varphi||\varphi') &= \int_{Support(\varphi)} \frac{d\varphi}{dm}\left[\ln \frac{d\varphi/dm}{d\varphi'/dm}\right] dm && (16) \\ &= E_\varphi\left[\ln \frac{d\varphi}{d\varphi'}\right]. && (17) \end{aligned}$$

[9]To see this rewrite, for any Θ', $\pi^t_{\mu_0}(\Theta')$ as conditional expectation of the indicator function of Θ' with respect to the induced probability measure P on the joint index and parameter space $(\Theta \times \Omega^\infty, \mathcal{F} \otimes \Sigma^\infty)$. To be precise, for the notation of our set-up it holds that, for all $\theta \in \Theta$ and $B \in \Sigma^\infty$,

$$P_\theta(B) \equiv P(B \mid \theta) \equiv P(\Theta \times B \mid \{\theta\} \times \Omega^\infty)$$

as well as, for all $\Theta' \in \mathcal{F}$,

$$\mu_0(\Theta') \equiv P(\Theta') \equiv P(\Theta' \times \Omega^\infty).$$

By Theorem 35.6 in Billingsley (1995) (which is an implication of the martingale convergence theorem), we obtain

$$\begin{aligned} \pi^t_{\mu_0}(\Theta') &\equiv E\left[I_{\Theta'}(\theta), P(\theta \mid X_1, ..., X_t)\right] \\ &\to E\left[I_{\Theta'}(\theta), P(\theta \mid X_1, X_2, ...)\right] \equiv \pi^\infty_{\mu_0}(\Theta') \end{aligned}$$

with P probability one.

[10]An accessible proof can be found in Section 1.3.3. of Gosh and Ramamoorthi (2003).

In our set-up, the KL-divergence (16) can take on values in $[0, \infty)$ whereby $D_{KL}(\varphi||\varphi') = 0$ if, and only if, $\varphi = \varphi'$ so that the KL-divergence gives us some notion about how close φ' is to φ without being a fully fledged metric.[11] In the continuous case, (16) becomes

$$D_{KL}(\varphi||\varphi') = \int_{x \in Support(\varphi)} \ln \left[\frac{d\varphi/dm}{d\varphi'/dm}(x) \right] \frac{d\varphi}{dm}(x)\, dx \tag{18}$$

where $\frac{d\varphi}{dm}$ and $\frac{d\varphi'}{dm}$ are the pdf's of the (absolutely continuous) distribution functions φ and φ', respectively. In the finite case, we have that

$$D_{KL}(\varphi||\varphi') = \sum_{\omega \in \Omega} d\varphi(\omega) \cdot \ln \frac{d\varphi(\omega)}{d\varphi'(\omega)} \tag{19}$$

where $d\varphi(\omega)$ and $d\varphi'(\omega)$ denote the probabilities of event $\{\omega\}$ with respect to the probability measures φ and φ', respectively.

Example 3. Revisit Example 1 and suppose that the φ_θ are given as normal distributions $\mathcal{N}(\mu_\theta, \sigma_\theta)$ with $\Theta = \{1, ..., n\}$. A straightforward exercise[12] shows that, for any $i, j \in \Theta$,

$$D_{KL}(\varphi_i||\varphi_j) = \ln \frac{\sigma_j}{\sigma_i} + \frac{\sigma_i^2 + (\mu_i - \mu_j)^2}{2\sigma_j^2} - \frac{1}{2}. \tag{20}$$

□

Theorem 0. (Berk 1966). The emerging posterior $\pi_{\mu_0}^\infty$ of a–not necessarily well-specified–prior μ_0 will almost surely concentrate at the subset $\Theta_{\mu_0}^* \subseteq Support(\mu_0)$ consisting of the KL-divergence minimizers φ_θ from the true measure φ_{θ^*}. That is,

$$Support\left(\pi_{\mu_0}^\infty\right) = \Theta_{\mu_0}^*, \text{ a.s. } P_{\theta^*} \tag{21}$$

such that

$$\Theta_{\mu_0}^* = \arg \min_{\theta \in Support(\mu_0)} D_{KL}(\varphi_{\theta^*}||\varphi_\theta). \tag{22}$$

[11] The KL-divergence is asymmetric and does not satisfy the triangle inequality. Note that, by convention, the KL-divergence takes on the value ∞ iff $d\varphi/dm > 0$ and $d\varphi'/dm = 0$; however, this case is not relevant to our paper because each φ' has full support on (Ω, Σ).

[12] For an elegant way to derive (20) see the answer of user 'ogrisel' under http://stats.stackexchange.com/questions/7440/kl-divergence-between-two-univariate-gaussians

Remark. While Berk (1966) does not explicitly mention the notion 'KL-divergence', which is implicit in his analysis by his Definition (b) of $\eta(\theta)$, Kleijn and Vaart (2006) do. To see that Berk's Theorem entails Doob's Theorem just observe that, for any well-specified prior μ_0,

$$\{\theta^*\} = \arg \min_{\theta \in Support(\mu_0)} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) \tag{23}$$

so that (21) becomes (15).

3 Bayesian learning with multiple priors

3.1 Prior selection rules

Turn now to a Bayesian decision maker who expresses ambiguity attitudes through multiple priors over the parameter values in Θ. Instead of an unique prior μ_0, we now consider a non-empty, closed set $\mathcal{M}_0 \subseteq \triangle^n$ of priors over Θ. Suppose, for the moment, that the decision maker forms posteriors from all his priors. Then he will, almost surely, end up with the following set of emerging posteriors after observing an unlimited amount of statistical information

$$\Pi^\infty = \bigcup_{\mu_0 \in \mathcal{M}_0} \left\{ \pi_{\mu_0}^\infty \right\}, \text{ a.s. } P_{\theta^*}. \tag{24}$$

If all priors in $\mathcal{M}_0$ are well-specified, we can immediately restate, by Doob's Theorem, M-2002's main finding according to which all emerging posteriors concentrate at the true measure φ_{θ^*}, i.e.,

$$\Pi^\infty = \{\delta_{\theta^*}\}, \text{ a.s. } P_{\theta^*}. \tag{25}$$

For the more general case of not-necessarily well-specified priors, we obtain, by Berk's Theorem 0, that every $\pi_{\mu_0}^\infty$ in (24) has support on $\Theta_{\mu_0}^*$ given by (22). In particular, if $\delta_\theta \in \mathcal{M}_0$, then also $\delta_\theta \in \Pi^\infty$. As a consequence, the set of emerging posteriors (24) might become quite implausible if not all priors are well-specified.

Example 4. Revisit the "Coin tossing" Example 2. Suppose that the parameter space is given as

$$\Theta = \{0.01, 0.99\} \tag{26}$$

such that $\theta^* = 0.99$. In the long run the decision maker will thus observe, by the law of large numbers, about 99% of all coin tosses resulting in Tails.

Further suppose that the set of priors is given as all probability measures on $(\Theta, \mathcal{F})$, i.e., $\mathcal{M}_0 = \triangle^2$. By Berk's Theorem 0, we obtain the following set of emerging posteriors

$$\Pi^\infty = \{\delta_{0.01}, \delta_{0.99}\} \tag{27}$$

because all mixed priors as well as $\delta_{0.99}$ converge to $\delta_{0.99}$ whereas $\delta_{0.01}$ converges to $\delta_{0.01}$. That is, after unlimited Bayesian learning the decision maker regards it still as possible that the objective probability of Tails might be only 1% despite having observed Tails in 99% of all coin tosses.□

On the one hand, we regard the set of emerging posteriors in the above example as highly unrealistic; i.e., we do not believe that there are many real-life decision makers who would end up with $\delta_{0.01} \in \Pi^\infty$. On the other hand, we regard it as too restrictive to consider only decision makers with well-specified priors as M-2002. In particular, we do not see any plausible reason why multiple priors decision makers should not hold some misspecified before they observe any data.

To resolve this "plausibility dilemma", we follow the seminal approach of ES-2007 and assume that the decision maker tests the plausibility of his priors against the observed data in accordance with some prior-selection rule. Formally, the set of admissible (=non-rejected) priors in the light of any observed data sample drawn from $X_1, ..., X_t$ is thereby determined by some prior-selection rule R, i.e.,

$$X_1, ..., X_t \longmapsto \mathcal{M}_{0,R}^t,\ t = 1, 2, ..., \tag{28}$$

such that the set of R-admissible priors at t, denoted $\mathcal{M}_{0R}^t$, satisfies, for all t,

$$\mathcal{M}_{0,R}^t \subseteq \mathcal{M}_0. \tag{29}$$

Note that, by (29), previously rejected priors might become admissible later on if they are supported by new data. To illustrate the concept of a prior-selection rule, consider the following two examples of perceivable rules.

Example 5. If the decision maker applies the *maximum expected likelihood rule*, he rejects each prior as implausible that does not maximize the expected likelihood for the observed data sample, i.e.,

$$\mathcal{M}_{0,ML}^t = \arg\max_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} \prod_{i=1}^{t} \frac{d\varphi_\theta}{dm} \cdot \mu_0(\theta). \tag{30}$$

□

Example 6. Now consider the *maximum expected loglikelihood rule.* By this rule, the decision maker rejects each prior as implausible that does not maximize the expected loglikelihood for the observed data sample, i.e.,

$$\mathcal{M}_{0,MLL}^{t}=\arg\max_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta}\ln\prod_{i=1}^{t}\frac{d\varphi_\theta}{dm}\cdot\mu_0(\theta)\,. \tag{31}$$

□

If $\mathcal{M}_0$ only contains degenerate measures, both rules are equivalent because likelihood and loglikelihood maximizers are identical. This equivalence is no longer the case if the expectation is taken with respect to non-degenerate priors in $\mathcal{M}_0$. More specifically, compared to the maximum expected likelihood rule the maximum expected loglikelihood rule punishes more strongly priors μ_0 that have θ's with small likelihoods in their support.

Remark. The above rules are very restrictive in that they (typically) reject all priors except for one. As a consequence, both rules generate a sequence of singleton sets of admissible priors to the effect that any ambiguity already vanishes after observing the first drawing, i.e., data-point. In the remainder of this paper, we therefore consider two families of less extreme prior-selection rules–the ES-2007 α-maximum expected likelihood, on the one hand, and the γ-maximum expected loglikelihood rule, on the other hand–which nest the maximum expected likelihood (resp. loglikelihood) rule as respective special cases.

3.2 Admissible limit priors and emerging posteriors

To describe the long-run learning behavior of a multiple priors decision maker who applies a prior-selection rule, we have to make a stand about how to define the set of admissible priors that survive this prior-selection rule if the number of data observations gets arbitrarily large. More precisely, we have to decide whether we either consider the *cluster* or the *limit points* of any sequence of R-admissible priors $\left\{\mathcal{M}_{0,R}^{t}\right\}_{t\in\mathbb{N}}$ as the priors in $\mathcal{M}_0$ that survive the prior-selection rule R in the limit.

Denote by $\overline{\lim}_{t\to\infty}\mathcal{M}_{0,R}^{t}$ the set that contains all *cluster points* in $\triangle^n$ of a given sequence $\left\{\mathcal{M}_{0,R}^{t}\right\}_{t\in\mathbb{N}}$. Formally, $\mu\in\triangle^n$ is a *cluster point of* $\left\{\mathcal{M}_{0,R}^{t}\right\}_{t\in\mathbb{N}}$ if, and only if, for every open set V around μ there are infinitely many t such that $V\cap\mathcal{M}_{0,R}^{t}\neq\emptyset$.[13]

[13]The set of all cluster points of a given sequence of sets is also called the *topological lim sup* of this sequence (Aliprantis and Border 2006, p. 114) or the *upper limit* of this sequence (Berge 1997, p. 119).

Conversely, denote by $\underline{\lim}_{t\to\infty}\mathcal{M}_{0,R}^{t}$ the set of all *limit points of the sequence* $\left\{\mathcal{M}_{0,R}^{t}\right\}_{t\in\mathbb{N}}$ such that $\mu\in\triangle^{n}$ is a *limit point of* $\left\{\mathcal{M}_{0,R}^{t}\right\}_{t\in\mathbb{N}}$ if, and only if, for every open set V around μ there exists some T such that, for all $t\geq T$, $V\cap\mathcal{M}_{0,R}^{t}\neq\emptyset$.[14] Whereas every limit point is a cluster point the converse is not true, implying

$$\underline{\lim}_{t\to\infty}\mathcal{M}_{0,R}^{t}\subseteq\overline{\lim}_{t\to\infty}\mathcal{M}_{0,R}^{t}. \tag{32}$$

To see the difference between both concepts of topological set limits applied to the notion of admissible limit priors consider the following example.

Example 7. Let φ_{θ} be the normal distribution with mean θ and variance 1 and suppose that $\Theta=\{\theta^{*},\theta_{1},\theta_{2}\}$ with $\theta^{*}=0,\theta_{1}=-1,\theta_{2}=1$. Further, suppose that $\mathcal{M}_{0}=\{\mu_{0}^{\prime},\mu_{0}^{\prime\prime}\}$ with

$$\mu_{0}^{\prime}=\delta_{\theta_{1}},\ \mu_{0}^{\prime\prime}=\delta_{\theta_{2}}. \tag{33}$$

That is, the decision maker will observe a sample that is generated by a symmetric (unbiased) random walk whereas he assumes that the data was either generated by a negatively or by a positively biased random walk. Further suppose that the decision maker applies the maximum expected likelihood rule[15] as prior-selection rule. Observe that

$$\sum_{\theta\in\Theta}\prod_{i=1}^{t}\frac{d\varphi_{\theta}}{dm}\cdot\mu_{0}^{\prime}(\theta)\ \geq\ \sum_{\theta\in\Theta}\prod_{i=1}^{t}\frac{d\varphi_{\theta}}{dm}\cdot\mu_{0}^{\prime\prime}(\theta)\Leftrightarrow \tag{34}$$

$$\prod_{k=1}^{t}\frac{d\phi_{\theta_{1}}}{dm}(X_{k})\ \geq\ \prod_{k=1}^{t}\frac{d\phi_{\theta_{2}}}{dm}(X_{k})\Leftrightarrow \tag{35}$$

$$\ln\frac{\prod_{k=1}^{t}\frac{d\phi_{\theta_{1}}}{dm}(X_{k})}{\prod_{k=1}^{t}\frac{d\phi_{\theta_{2}}}{dm}(X_{k})}\ \geq\ 0\Leftrightarrow \tag{36}$$

$$\ln\frac{\exp\left[-\frac{1}{2}\sum_{k=1}^{t}(X_{k}-\theta_{1})^{2}\right]}{\exp\left[-\frac{1}{2}\sum_{k=1}^{t}(X_{k}-\theta_{2})^{2}\right]}\ \geq\ 0\Leftrightarrow \tag{37}$$

$$\ln\exp\left[(\theta_{1}-\theta_{2})\sum_{k=1}^{t}X_{k}-\frac{t}{2}\left(\theta_{1}^{2}-\theta_{2}^{2}\right)\right]\ \geq\ 0\Leftrightarrow \tag{38}$$

$$-2\sum_{k=1}^{t}X_{k}\ \geq\ 0 \tag{39}$$

[14] The set of all limit points of a given sequence of sets is also called the *topological lim inf* of this sequence (Aliprantis and Border 2006, p. 114) or the *lower limit* of this sequence (Berge 1997, p. 119).

[15] Which is here, due to the degenerate priors, equivalent to the maximum expected loglikelihood rule.

implying that

$$\mathcal{M}_{0,ML}^{t}=\begin{cases}\{\mu_0'\} & \text{if } \sum_{k=1}^{t}X_k<0\\ \{\mu_0',\mu_0''\} & \text{if } \sum_{k=1}^{t}X_k=0\\ \{\mu_0''\} & \text{if } \sum_{k=1}^{t}X_k>0\end{cases} \tag{40}$$

By the recurrence theorem (Chung and Fuchs 1951), we will almost surely observe that $\sum_{k=1}^{t} X_k$ crosses the zero line infinitely many times if t gets arbitrarily large. Consequently, there do not exist any limit points for the sequence $\left\{\mathcal{M}_{0,ML}^{t}\right\}_{t\in\mathbb{N}}$ so that

$$\underline{\lim}_{t\to\infty}\mathcal{M}_{0,ML}^{t}=\emptyset,\text{ a.s. } P_{\theta^*}. \tag{41}$$

On the other hand, we obtain the non-empty set of cluster points

$$\overline{\lim}_{t\to\infty}\mathcal{M}_{0,ML}^{t}=\{\mu_0',\mu_0''\}\,,\text{ a.s. } P_{\theta^*}. \tag{42}$$

□

To take $\overline{\lim}_{t\to\infty}\mathcal{M}_{0,ML}^{t}=\{\mu_0',\mu_0''\}$ rather than the empty set $\underline{\lim}_{t\to\infty}\mathcal{M}_{0,ML}^{t}$ as the set of one-admissible limit priors in the above example appears to us as the natural thing to do. Since μ_0' as well as μ_0'' will always be supported (almost surely) by some data, a cautious (conservative) decision maker should not rule out any prior in $\{\mu_0',\mu_0''\}$ as impossible. Motivated by these considerations, we introduce the following definition of R-admissible limit priors.

Definition 2. We define set of *R-admissible limit priors*, denoted $\mathcal{M}_{0,R}^{\infty}$, as the set of cluster points of $\left\{\mathcal{M}_{0,R}^{t}\right\}_{t\in\mathbb{N}}$ that almost surely emerge, i.e.,

$$\mathcal{M}_{0,R}^{\infty}\equiv\overline{\lim}_{t\to\infty}\mathcal{M}_{0,R}^{t},\text{ a.s. } P_{\theta^*}. \tag{43}$$

In words: The set of R-admissible limit priors consists of all priors in $\mathcal{M}_0$ that will almost surely pop up again as elements in sets of the sequence $\left\{\mathcal{M}_{0,R}^{t}\right\}_{t\in\mathbb{N}}$ whenever the data sample becomes arbitrarily large. In a next step, we assume that, for any given prior-selection rule R, all emerging posteriors must have been formed from R-admissible limit priors.[16]

[16] Equivalently, we define the emerging posteriors as the cluster points of the sequence of sets of posteriors $\{\Pi_R^t\}_{t\in\mathbb{N}}$ that almost surely emerge.

Definition 3. The set of *emerging posteriors* under the prior-selection rule R, denoted Π_R^∞, is defined as

$$\Pi_R^\infty = \bigcup_{\mu_0 \in \mathcal{M}_{0,R}^\infty} \left\{ \pi_{\mu_0}^\infty \right\}. \tag{44}$$

We say that *ambiguity vanishes* if, and only if, Π_R^∞ is a singleton, i.e.,

$$\Pi_R^\infty = \left\{ \pi_{\mu_0}^\infty \right\}. \tag{45}$$

Suppose, for example, that every prior in $\mathcal{M}_{0,R}^\infty$ has a unique KL-divergence minimizer in its support. Then (44) becomes, by Berk's Theorem, the following collection of Dirac measures

$$\Pi_R^\infty = \bigcup_{\mu_0 \in \mathcal{M}_{0,R}^\infty} \left\{ \delta_{\hat{\theta}} \mid \hat{\theta} \in \arg \min_{\theta \in Support(\mu_0)} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) \right\}. \tag{46}$$

In that case, vanishing ambiguity means

$$\Pi_R^\infty = \{\delta_{\hat{\theta}}\} \tag{47}$$

for some $\hat{\theta} \in \Theta$ whereby we allow for the possibility that $\hat{\theta} \neq \theta^*$. That is, vanishing ambiguity does not necessarily imply that the decision maker also learns the truth.

3.3 The Epstein and Schneider (2007) α-maximum expected likelihood rule

The α-maximum expected likelihood rule, introduced by ES-2007, relaxes the maximum expected likelihood rule of Example 3 by allowing the decision maker to keep priors that are α-close to the expected likelihood maximizing prior. Restricted to the one-urn environment, the formal definition of this rule is given as follows.

Definition 4. The *α-maximum expected likelihood rule* (ES-2007). Fix some $\alpha \in (0, 1]$. The set of *α-admissible priors* after observing a sample drawn from $X_1, ..., X_t$ is given as

$$\mathcal{M}_{0,\alpha}^t = \left\{ \mu_0' \in \mathcal{M}_0 \mid \sum_{\theta \in \Theta} \prod_{i=1}^t \frac{d\varphi_\theta}{dm} \cdot \mu_0' (\theta) \geq \alpha \cdot \max_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} \prod_{i=1}^t \frac{d\varphi_\theta}{dm} \cdot \mu_0 (\theta) \right\}. \tag{48}$$

Under some regularity assumptions[17], ES-2007 derive their Claim 3 (p. 1301) according to which all α-admissible limit priors will be well-specified if there exists at least one well-specified prior in $\mathcal{M}_0$.

Claim 3 in ES-2007. If $\mu_0(\theta^*) > 0$ for some $\mu_0 \in \mathcal{M}_0$, then $\mu_0(\theta^*) > 0$ for all $\mu_0 \in \mathcal{M}_{0,\alpha}^{\infty}$.

To see that this result by ES-2007 is surprisingly strong, consider the following example.

Example 8. Revisit the coin tossing situation of Example 2 and suppose that the parameter space is given as

$$\Theta = \{0.49, 0.5, 0.99\} \tag{49}$$

where $\theta \in \Theta$ is the probability of event $\{Tails\}$. Further suppose that the coin is slightly unfair such that $\theta^* = 0.49$. Next consider the set of priors $\mathcal{M}_0 = \{\mu_0', \mu_0''\}$ such that

$$\mu_0' = \varepsilon \cdot \delta_{0.49} + (1 - \varepsilon) \cdot \delta_{0.99}, \tag{50}$$

$$\mu_0'' = \delta_{0.5}, \tag{51}$$

for some small $\varepsilon > 0$. Note that μ_0' is well-specified but a very incorrect belief because it attaches the large probability $1 - \varepsilon$ to the very false parameter value $\theta = 0.99$. On the other hand, μ_0'' is misspecified but very close to the true value. If the decision maker applies the α-maximum expected likelihood rule, he will, by Claim 3 in ES-2007, eventually reject the prior μ_0'' so that μ_0' remains the only admissible prior from which he forms posteriors. That is, the α-maximum expected likelihood rule drives out the almost true parameter value 0.5 in favor of the prior $\varepsilon \cdot \delta_{0.49} + (1 - \varepsilon) \cdot \delta_{0.99}$, which–as a belief–is quite off-the-mark.□

Based on Claim 3 in ES-2007, we can immediately derive, by an application of Doob's Theorem, Epstein and Schneider's (2007) Theorem 1 (p. 1288) for our one-urn environment as follows.

[17] ES-2007 restrict attention to a finite state space Ω. Because ES-2007 admit for non-finite index sets Θ, they impose weak compactness of $\mathcal{M}_0$ and they also require that $\mu_0(\theta^*)$ has to be uniformly bounded away from zero if θ^* is in the support of μ_0. For our finite index sets, $\mathcal{M}_0$ is weakly compact if, and only if, it is closed whereby the bounded-away-from-zero condition is automatically satisfied for finite index sets. For further details about their regularity assumptions see Theorem 1 (ES-2007, p. 1288).

Theorem 1 (ES-2007). Suppose that Ω is finite. If $\mu_0(\theta^*) > 0$ for some $\mu_0 \in \mathcal{M}_0$, then the set of posteriors that emerge under the α-maximum expected likelihood rule is given as

$$\Pi_\alpha^\infty = \{\delta_{\theta^*}\}. \tag{52}$$

4 New results: The γ-maximum expected loglikelihood rule

Central to our paper is the introduction of a new prior-selection rule as a perceivable alternative to the ES-2007 α-maximum expected likelihood rule.

Definition 5. The *γ-maximum expected loglikelihood rule.* Fix some $\gamma \in [1, \infty)$. The set of admissible priors after observing a sample drawn from $X_1, ..., X_t$ is given as

$$\mathcal{M}_{0,\gamma}^t = \left\{ \mu_0' \in \mathcal{M}_0 \mid \sum_{\theta \in \Theta} \ln \prod_{i=1}^t \frac{d\varphi_\theta}{dm} \cdot \mu_0'(\theta) \geq \gamma \cdot \max_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} \ln \prod_{i=1}^t \frac{d\varphi_\theta}{dm} \cdot \mu_0(\theta) \right\} \tag{53}$$

whenever the maximal expected loglikelihood is not strictly positive, i.e.,

$$\max_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} \ln \prod_{i=1}^t \frac{d\varphi_\theta}{dm} \cdot \mu_0(\theta) \leq 0, \tag{54}$$

and

$$\mathcal{M}_{0,\gamma}^t = \mathcal{M}_{0,MLL}^t = \arg \max_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} \ln \prod_{i=1}^t \frac{d\varphi_\theta}{dm} \cdot \mu_0(\theta), \tag{55}$$

else.

If (54) holds, the decision maker judges, by (53), priors as plausible whose expected loglikelihood is γ-close to the expected loglikelihood of the maximizing prior. Note that the greater the value of $\gamma \geq 1$, the more priors will be included in $\mathcal{M}_{0,\gamma}^t$ for a given data sample. Further note that (54) always holds for the finite but not necessarily for the continuous case (e.g., let $\arg\max_{\mu_0 \in \mathcal{M}_0} = \delta_\theta$ such that φ_θ is the uniform distribution on $[a, b]$ with $0 < a, b < 1$). If (54) is violated to the effect that the decision maker deals with a strictly positive maximal expected loglikelihood, we simply assume that the γ-maximum expected loglikelihood rule reduces to the maximum expected loglikelihood rule (55).

Proposition 1. The set of posteriors that emerge under the γ-maximum expected loglikelihood rule is given as

$$\Pi_{\gamma}^{\infty} = \bigcup_{\mu_0 \in \mathcal{M}_{0,\gamma}^{\infty}} \left\{ \pi_{\mu_0}^{\infty} \right\} \tag{56}$$

such that either

$$\mathcal{M}_{0,\gamma}^{\infty} = \left\{ \mu_0' \in \mathcal{M}_0 \mid \sum_{\theta \in \Theta} D_{KL}(\varphi_{\theta^*} || \varphi_{\theta}) \cdot \mu_0'(\theta) \leq \min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} D_{KL}(\varphi_{\theta^*} || \varphi_{\theta}) \cdot \mu_0(\theta) \right.$$
$$\left. + \frac{(1-\gamma)}{\gamma} \sum_{\theta \in \Theta} E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_{\theta}}{dm} \right] \cdot \mu_0'(\theta) \right\} \tag{57}$$

whenever this set is non-empty, or

$$\mathcal{M}_{0,\gamma}^{\infty} = \arg \min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} D_{KL}(\varphi_{\theta^*} || \varphi_{\theta}) \cdot \mu_0(\theta) \tag{58}$$

else.[18]

Observe that (57) is empty if, and only if, the *expected cross-entropy*

$$-\sum_{\theta \in \Theta} E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_{\theta}}{dm} \right] \cdot \mu_0'(\theta) \tag{59}$$

is strictly negative and $\gamma > 1$. A strictly negative expected cross-entropy is impossible for the finite but not for the continuous case. If the expected cross-entropy is positive, i.e., (57) is non-empty, the emerging posteriors are formed from priors whose expected KL-divergence is sufficiently close to the minimal expected KL-divergence whereby greater values of γ will imply greater sets of posteriors. The following subsections further characterize the set of emerging posteriors (56) for different values of γ.

4.1 Special case $\gamma = 1$: The maximum expected loglikelihood rule

For $\gamma = 1$ the γ-maximum expected loglikelihood rule becomes the maximum expected loglikelihood rule of Example 6, i.e., (57) becomes

$$\mathcal{M}_{0,\gamma=1}^{\infty} = \arg \min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} D_{KL}(\varphi_{\theta^*} || \varphi_{\theta}) \cdot \mu_0(\theta) \, . \tag{60}$$

[18] The proof of Proposition 1 uses our assumption of a finite index space Θ. We conjecture that an analogous proof should (by the dominated convergence theorem) also go through for an infinite state space Θ combined with a finite Ω (as in ES-2007). However, the situation might be different for the continuous case if the pdf's are not bounded away from zero.

That is, under the maximum expected loglikelihood rule, the admissible limit priors are the priors that minimize the expected KL-divergence from the true measure.

Corollary 1. Suppose that $\gamma = 1$. If $\mu_0' (\theta^*) > 0$ for some

$$\mu_0' \in \arg \min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) \cdot \mu_0 (\theta), \tag{61}$$

then the set of posteriors that emerge under the maximum expected loglikelihood rule is given as

$$\Pi_{\gamma=1}^{\infty} = \{\delta_{\theta^*}\}. \tag{62}$$

In particular, (62) holds if $\delta_{\theta^*} \in \mathcal{M}_0$.

In contrast to the mere existence of a well specified prior in Theorem 1 for the ES-2007 model, the well-specified prior μ_0' of Corollary 1 must also minimize the expected KL-divergence. The following example demonstrates that our decision maker does not necessarily learns the truth if there is some well-specified prior but $\delta_{\theta^*} \notin \mathcal{M}_0$.

Example 9. Revisit the coin tossing Example 8 where

$$\Theta = \{0.49, 0.5, 0.99\} \tag{63}$$

with $\theta^* = 0.49$ and $\mathcal{M}_0 = \{\mu_0', \mu_0''\}$ such that

$$\begin{aligned} \mu_0' &= \varepsilon \cdot \delta_{0.49} + (1-\varepsilon) \cdot \delta_{0.99}, && (64) \\ \mu_0'' &= \delta_{0.5}. && (65) \end{aligned}$$

By continuity of the KL-divergence, we can always find $\varepsilon > 0$ sufficiently small such that

$$\begin{aligned} \sum_{\theta \in \Theta} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) \cdot \mu_0'' (\theta) &< \sum_{\theta \in \Theta} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) \cdot \mu_0' (\theta) && (66) \\ &\Leftrightarrow \\ E_{\varphi_{\theta^*}} [\ln d\varphi_{\theta^*}] - E_{\varphi_{\theta^*}} [\ln d\varphi_{0.5}] &< \varepsilon \cdot \left(E_{\varphi_{\theta^*}} [\ln d\varphi_{\theta^*}] - E_{\varphi_{\theta^*}} [\ln d\varphi_{\theta^*}]\right) && (67) \\ &\quad + (1-\varepsilon) \cdot \left(E_{\varphi_{\theta^*}} [\ln d\varphi_{\theta^*}] - E_{\varphi_{\theta^*}} [\ln d\varphi_{0.99}]\right) \\ &\Leftrightarrow \\ E_{\varphi_{\theta^*}} [\ln d\varphi_{0.5}] &> \varepsilon \cdot E_{\varphi_{\theta^*}} [\ln d\varphi_{\theta^*}] + (1-\varepsilon) \cdot E_{\varphi_{\theta^*}} [\ln d\varphi_{0.99}] \end{aligned}$$

since

$$E_{\varphi_{\theta^*}}\left[\ln d\varphi_{\theta^*}\right] > E_{\varphi_{\theta^*}}\left[\ln d\varphi_{0.5}\right] > E_{\varphi_{\theta^*}}\left[\ln d\varphi_{0.99}\right]. \tag{68}$$

To be concrete observe that

$$\begin{aligned} E_{\varphi_{\theta^*}}\left[\ln d\varphi_{0.5}\right] &= (1-0.49)\cdot\ln(1-0.5)+0.49\cdot\ln 0.5 \approx -0.693147 \\ E_{\varphi_{\theta^*}}\left[\ln d\varphi_{\theta^*}\right] &= (1-0.49)\cdot\ln(1-0.49)+0.49\cdot\ln 0.49 \approx -0.692947 \\ E_{\varphi_{\theta^*}}\left[\ln d\varphi_{0.9999}\right] &= (1-0.49)\cdot\ln(1-0.99)+0.49\cdot\ln 0.99 \approx -2.35356 \end{aligned}$$

so that any $\varepsilon > 0$ satisfying

$$\varepsilon < \frac{E_{\varphi_{\theta^*}}\left[\ln d\varphi_{0.5}\right]-E_{\varphi_{\theta^*}}\left[\ln d\varphi_{0.99}\right]}{E_{\varphi_{\theta^*}}\left[\ln d\varphi_{\theta^*}\right]-E_{\varphi_{\theta^*}}\left[\ln d\varphi_{0.99}\right]} \tag{69}$$

$$\Leftrightarrow$$

$$\varepsilon < 0.99988 \tag{70}$$

would do. □

Note that we obtain in the above example for $\varepsilon < 0.99988$ that

$$\begin{aligned} \Pi^{\infty} &= \{\delta_{0.49}, \delta_{0.5}\}, \\ \Pi^{\infty}_{\alpha} &= \{\delta_{0.49}\}, \\ \Pi^{\infty}_{\gamma=1} &= \{\delta_{0.5}\}. \end{aligned}$$

Without any prior-selection rule ambiguity will not vanish. Under the ES-2007 α-maximum expected likelihood rule ambiguity will, for all $\alpha > 0$, vanish whereby the decision maker learns the true probability measure $\varphi_{0.49}$. Ambiguity will also vanish under the maximum expected loglikelihood rule, however, here the decision maker will learn the almost true probability measure $\varphi_{0.5}$ rather than the true $\varphi_{0.49}$.

Remark. To see the intuition behind the formal difference between the maximal expected likelihood versus the maximal expected loglikelihood rule, consider the analogy to risk-neutral versus strictly risk averse EU maximization with respect to multiple priors. If likelihoods are taken as prizes, expected likelihood maximization corresponds to risk neutral expected utility (=expected value) maximization. In contrast, expected loglikelihood maximization corresponds to strictly risk averse expected utility maximization such that the utils are given as the logs of the prizes. By this interpretation, a decision maker who uses loglikelihoods as utils is more cautious (risk-averse) than a decision maker who instead uses likelihoods as utils. In particular, priors that put positive weight on likelihoods that are close to zero will be considered as highly unfavorable from the perspective of such a cautious decision maker. In the above example, the well-specified prior $\varepsilon \cdot \delta_{0.49} + (1-\varepsilon) \cdot \delta_{0.99}$ is rejected as implausible because the positive weight on the unlikely parameter $\theta = 0.99$ pulls down the expected loglikelihood of this prior.

4.2 Allowing for sufficiently large $\gamma > 1$

Turn now to the case that $\gamma > 1$. The following result implies that the emerging posteriors in (56) may concentrate at different indices whereby there tend to be more emerging posteriors for greater values of γ.

Proposition 2. Suppose that φ_{θ^*} satisfies

$$E_{\varphi_{\theta^*}}\left[\ln \frac{d\varphi_{\theta^*}}{dm}\right] < 0. \tag{71}$$

Then

$$\delta_{\hat{\theta}} \in \Pi_{\gamma}^{\infty} \tag{72}$$

for some sufficiently large γ if, and only if,

$$\left\{\hat{\theta}\right\} = \arg \min_{\theta \in Support\left(\mu_0'\right)} D_{KL}(\varphi_{\theta^*} || \varphi_{\theta}) \tag{73}$$

for some prior $\mu_0' \in \mathcal{M}_0$.

Note that the condition (71) is equivalent to the condition that the *entropy*[19] of φ_{θ^*} has to be strictly positive. Condition (71) is thus always satisfied for the discrete case since φ_{θ^*} has, by assumption, full support on Ω with $\#\Omega > 1$. The situation is different for the continuous case where pdf's can take on values greater than one so that (71) might become positive.

Corollary 2. Suppose that (71) holds. For any given $\theta \in \Theta$, if $\delta_\theta \in \mathcal{M}_0$, then

$$\delta_{\theta} \in \Pi_{\gamma}^{\infty} \tag{74}$$

for some sufficiently large γ.

[19] More precisely, we refer to the *Shannon entropy* for the discrete and to the *differential entropy* for the continuous case.

In contrast to the α-maximum expected likelihood and to the maximum expected loglikelihood rule (i.e., $\gamma = 1$), ambiguity will thus not vanish under the γ-maximum expected loglikelihood rule for sufficiently large values of γ. We come back to this observation in the next section where we investigate non-vanishing STP-violations.

5 Application: Non-vanishing violations of the sure-thing principle

This section embeds Ellsberg's (1961) original one-urn experiment within our model of Bayesian learning with multiple priors under the assumption that the decision maker is a *maxmin expected utility* decision maker. Two main findings emerge. First, we can establish the possibility of non-vanishing violations of Savage's (1954) sure-thing principle. Second, we demonstrate that the non-vanishing ambiguity increases in the γ parameter.

5.1 The Ellsberg one-urn experiment

Savage (1954) considers a decision maker who has preferences $\succeq$ over Savage acts which map some state space Ω into a set of consequences, denoted Z. By imposing several structural and behavioral axioms, Savage derives the celebrated *subjective expected utility* (SEU) representation of $\succeq$ such that, for all Savage acts f, g,

$$f \succeq g \Leftrightarrow \int_{\omega \in \Omega} u(f(\omega))\, d\varphi \geq \int_{\omega \in \Omega} u(g(\omega))\, d\varphi, \tag{75}$$

where the subjective probability measure φ as well as the utility function $u : Z \to \mathbb{R}$ are uniquely[20] pinned down by the decision maker's preferences. We introduce the following notational convention for the SEU of act f with respect to probability measure φ

$$EU(f, \varphi) \equiv \int_{\omega \in \Omega} u(f(\omega))\, d\varphi. \tag{76}$$

Savage's key behavioral axiom is the *sure-thing principle* which states that, for all Savage acts f, g, h and events $E \in \Sigma$,

$$f_E h \succeq g_E h \;\;\Leftrightarrow\;\; f_E h' \succeq g_E h' \tag{77}$$

whereby

$$f_E h(\omega) = \begin{cases} f(\omega) & \text{for } \omega \in E \\ h(\omega) & \text{for } \omega \in \neg E \end{cases} \tag{78}$$

[20] Of course, the utility function u is only unique up to some positive affine transformation.

Starting with Ellsberg (1961), several experiments have reported systematic violations of the sure-thing principle, also dubbed 'Ellsberg paradoxes'. Let us focus on Ellsberg's (1961, p. 654) original one-urn experiment. The Ellsberg urn contains 30 red balls and 60 black or yellow balls of unknown proportion. Define the relevant state space

$$\Omega = \{\omega_1, \omega_2, \omega_3\} \tag{79}$$

where ω_1 (resp. ω_2, ω_3) stands for the state in which a red (resp. black, yellow) ball will be drawn. Next consider the following four Savage acts where $E = \{\omega_1, \omega_2\}$

	ω_1	ω_2	ω_3
$f_E h$	1	0	0
$g_E h$	0	1	0
$f_E h'$	1	0	1
$g_E h'$	0	1	1

The majority of decision makers express the preferences

$$f_E h \succ g_E h \text{ and } g_E h' \succ f_E h'. \tag{80}$$

Note that the 'Ellsberg paradox' (80) constitutes a violation of the sure-thing principle (77) and can therefore not be accommodated by SEU theory.

5.2 Maxmin expected utility

To accommodate the Ellsberg paradox (80), Gilboa and Schmeidler (1989) propose a "maxmin expected utility with non-unique prior" (=MEU) representation such that, for all Savage acts f, g,

$$f \succeq g \Leftrightarrow \min_{\varphi \in \mathcal{P}} \int_{\omega \in \Omega} u(f(\omega))\, d\varphi \geq \min_{\varphi \in \mathcal{P}} \int_{\omega \in \Omega} u(g(\omega))\, d\varphi \tag{81}$$

for some non-empty set of probability measures $\mathcal{P}$.[21] If $\mathcal{P}$ reduces to a singleton, i.e., $\mathcal{P} = \{\varphi\}$ for any subjective probability measure φ, MEU reduces to SEU. However, if $\mathcal{P}$ does not reduce to a singleton, the decision maker's preferences express ambiguity in the sense that he cannot pin down his uncertainty through a unique probability measure.

[21] Gilboa and Schmeidler (1989) axiomatize MEU within an Anscombe-Aumann (1963) framework where the set of consequences Z contains all lotteries over some non-degenerate set of deterministic prizes. Under this Gilboa and Schmeidler (1989) axiomatization, $\mathcal{P}$ is uniquely pinned down as a non-empty, closed and convex set of finitely additive probability measures. We ignore here this specific axiomatic foundation and also allow for, e.g., non-convex $\mathcal{P}$.

The MEU concept assumes that ambiguity is always resolved in a very pessimistic way: Each act is evaluated with respect to the probability measure in $\mathcal{P}$ that gives the minimal expected utility for the act in question.

Although this assumption of extreme *ambiguity aversion* is, in general, somewhat unrealistic[22], we follow the majority of the literature and suppose that the preferences (80) can be best explained through extreme ambiguity aversion. That is, in the remainder of this section we consider a MEU decision maker and we denote by

$$MEU\left(f,\mathcal{P}\right)\equiv\min_{\varphi\in\mathcal{P}}EU\left(f,\varphi\right) \tag{82}$$

the decision maker's maxmin expected utility from act f with respect to the set of probability measures $\mathcal{P}$.

To see that MEU can indeed accommodate the Ellsberg paradox (80) let $\Omega=\{\omega_1,\omega_2,\omega_3\}$ and $\Sigma=2^{\Omega}$. Furthermore, (quite naturally) suppose that all probability measures φ in $\mathcal{P}$ have the following structure

$$\varphi=\left(\frac{1}{3},\varphi\left(\omega_2\right),\frac{2}{3}-\varphi\left(\omega_2\right)\right) \tag{83}$$

for some $\varphi\left(\omega_2\right)\in\left[0,\frac{2}{3}\right]$. Without loss of generality, set $u\left(z\right)=z$ for $z\in\{0,1\}$. By (83),

$$MEU\left(f_Eh,\mathcal{P}\right)=\frac{1}{3}\text{ and }MEU\left(g_Eh',\mathcal{P}\right)=\frac{2}{3}. \tag{84}$$

If $\varphi_1\in\mathcal{P}$ such that $\varphi_1\left(\omega_2\right)<\frac{1}{3}$, then

$$MEU\left(g_Eh,\mathcal{P}\right)\leq EU\left(g_Eh,\varphi_1\right)<\frac{1}{3}. \tag{85}$$

If $\varphi_2\in\mathcal{P}$ such that $\varphi_2\left(\omega_2\right)>\frac{1}{3}$, then

$$MEU\left(f_Eh',\mathcal{P}\right)\leq EU\left(f_Eh',\varphi_2\right)<\frac{2}{3}. \tag{86}$$

Collecting the above arguments gives us the following result.

Lemma 1. If there are $\varphi_1,\varphi_2\in\mathcal{P}$ such that

$$\varphi_1\left(\omega_2\right)<\frac{1}{3}\text{ and }\varphi_2\left(\omega_2\right)>\frac{1}{3}, \tag{87}$$

then the MEU decision maker commits the Ellsberg paradox (80), i.e.,

$$MEU\left(f_Eh,\mathcal{P}\right)>MEU\left(g_Eh,\mathcal{P}\right), \tag{88}$$

$$MEU\left(g_Eh',\mathcal{P}\right)>MEU\left(f_Eh',\mathcal{P}\right). \tag{89}$$

[22]For a more realistic generalization of MEU, see the α-MEU concept of Ghirardato et al. (2004).

5.3 Bayesian learning with multiple priors and non-vanishing STP violations

We recast the Ellsberg one-urn experiment within our formal set-up of Bayesian learning with multiple priors. Let $\Omega = \{\omega_1, \omega_2, \omega_3\}$ and $\Sigma = 2^\Omega$ and consider the following set of probability measures on (Ω, Σ)

$$\Phi = \left\{ \varphi_\theta = \left(\frac{1}{3}, \varphi_\theta(\omega_2), \frac{2}{3} - \varphi_\theta(\omega_2) \right) \mid \varphi_\theta(\omega_2) = \frac{\theta}{90}, \theta \in \Theta \right\} \tag{90}$$

such that the index set Θ is given as

$$\Theta = \{1, ..., 59\}. \tag{91}$$

The indices in Θ correspond to the possible numbers of black balls in the urn whereby we assume that there is at least one black (and one yellow) ball in the urn.[23] As a consequence, Φ contains all probability measures with full support on Ω that might be deemed possible by the decision maker if he associates probabilities with the possible ratios of balls in the Ellsberg urn. To focus our analysis, we assume that $\theta^* = 30$, i.e., we set $\varphi_{\theta^*}(\omega_2) = \frac{1}{3}$ as the true probability that a black ball will be drawn from the urn.

Consider at first the *a priori* decision situation in which the decision maker has not yet received any statistical information about θ^* in the form of θ^*-i.i.d. drawings. We model the MEU decision maker's *a priori* uncertainty about the true parameter value through some set of priors $\mathcal{M}_0$ defined on the index space $(\Theta, \mathcal{F})$. The set of probability measures $\mathcal{P}$ in (82) then becomes the set of reduced compound measures in $\mathcal{M}_0 \circ \Phi$ defined on (Ω, Σ) such that, for any Savage act f,

$$\begin{aligned} MEU(f, \mathcal{M}_0 \circ \Phi) &= \min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} \left[\sum_{\omega \in \Omega} u(f(\omega)) \varphi_\theta(\omega) \right] \mu_0(\theta) && (92) \\ &= \min_{\varphi^0 \in \mathcal{M}_0 \circ \Phi} \sum_{\omega \in \Omega} u(f(\omega)) \varphi^0(\omega) && (93) \end{aligned}$$

where

$$\varphi^0(\omega) \equiv \sum_{\theta \in \Theta} \varphi_\theta(\omega) \mu_0(\theta). \tag{94}$$

Now consider the *a posteriori* decision situation in which the decision maker had started out with priors in $\mathcal{M}_0$ and subsequently observed arbitrarily many θ^*-i.i.d. drawings. The MEU decision maker will (a.s. P_{θ^*}) resolve his uncertainty about the true

[23] We exclude $\theta = 0$ and $\theta = 60$ out of convenience since we do not want to make a stand about Bayesian updating in the light of events that the decision maker perceives as impossible. E.g., we want to avoid the case that $\mu_0 = \delta_{60}$ and the decision maker observes a yellow ball drawn from the urn.

parameter value through the set of emerging posteriors Π_γ^∞ defined on the index space $(\Theta, \mathcal{F})$. In this *a posteriori* decision situation, we thus obtain the set of reduced compound measures $\mathcal{P} = \Pi_R^\infty \circ \Phi$, defined on (Ω, Σ), implying, for any Savage act f,

$$\begin{aligned} MEU\left(f, \Pi_\gamma^\infty \circ \Phi\right) &= \min_{\left\{\pi_{\mu_0}^\infty \mid \mu_0 \in \mathcal{M}_{0,\gamma}^\infty\right\}} \sum_{\theta\in\Theta} \left[\sum_{\omega\in\Omega} u\left(f\left(\omega\right)\right) \varphi_\theta\left(\omega\right)\right] \pi_{\mu_0}^\infty\left(\theta\right) && (95) \\ &= \min_{\varphi_\gamma^\infty \in \Pi_\gamma^\infty \circ \Phi} \sum_{\omega\in\Omega} u\left(f\left(\omega\right)\right) \varphi_\gamma^\infty\left(\omega\right) && (96) \end{aligned}$$

where

$$\varphi_\gamma^\infty\left(\omega\right) \equiv \sum_{\theta\in\Theta} \varphi_\theta\left(\omega\right) \pi_{\mu_0}^\infty\left(\theta\right). \qquad (97)$$

Note that an *a priori* MEU decision maker becomes an *a posteriori* SEU decision maker through Bayesian learning if, and only if, the set of emerging posteriors Π_γ^∞ is a singleton. Since Π_γ^∞ only contains degenerate probability measures, such an *a posteriori* SEU decision maker will hold a subjective belief that coincides with the objective probability measure φ_{θ^*} on (Ω, Σ) if, and only if, $\Pi_\gamma^\infty = \{\delta_{\theta^*}\}$.

We speak of non-vanishing STP violations if the decision maker commits the Ellsberg paradox (80) in the *a priori* as well as in the *a posteriori* decision situation. More precisely, we say that STP violations do not vanish if, and only if,

$$\begin{aligned} MEU\left(f_E h, \mathcal{M}_0 \circ \Phi\right) &> MEU\left(g_E h, \mathcal{M}_0 \circ \Phi\right) \text{ and} && (98) \\ MEU\left(g_E h', \mathcal{M}_0 \circ \Phi\right) &> MEU\left(f_E h', \mathcal{M}_0 \circ \Phi\right) && (99) \end{aligned}$$

as well as

$$\begin{aligned} MEU\left(f_E h, \Pi_\gamma^\infty \circ \Phi\right) &> MEU\left(g_E h, \Pi_\gamma^\infty \circ \Phi\right) \text{ and} && (100) \\ MEU\left(g_E h', \Pi_\gamma^\infty \circ \Phi\right) &> MEU\left(f_E h', \Pi_\gamma^\infty \circ \Phi\right). && (101) \end{aligned}$$

Proposition 3. Suppose that there are two misspecified priors $\mu_0', \mu_0'' \in \mathcal{M}_0$ such that μ_0' has support only on indices $\theta < 30$ and μ_0'' has support only on indices $\theta > 30$. Then there exists some sufficiently large $\gamma < \infty$ such that STP violations do not vanish.

Remark. The reader should be careful to distinguish between the two different notions of 'multiple priors' used in the different strands of literature that are relevant to our paper. First, there is our notion of the set of priors $\mathcal{M}_0$, defined on $(\Theta, \mathcal{F})$, which captures the decision maker's (initial/unconditional) uncertainty about the true measure in Φ. This 'multiple priors' notion is in line with the literature on Bayesian learning/updating. Second, there is the set $\mathcal{P}$, defined on (Ω, Σ), of additive probability measures that appears in the MEU utility representation (82). The (axiomatic) decision theoretic literature typically refers to the members in $\mathcal{P}$ as the 'multiple priors' that are relevant to the agent's decision situation. With respect to this decision theoretic notion, our set of reduced compound measures $\mathcal{M}_0 \circ \Phi$ captures the 'multiple priors' relevant to the *a priori* decision situation whereas $\mathcal{M}_0 \circ \Pi_\gamma^\infty$ captures the 'multiple priors' relevant to the *a posteriori* decision situation.

5.4 The set of emerging posteriors is γ sensible

Ambiguity typically decreases but not necessarily vanishes through statistical learning in our model of Bayesian learning with multiple priors. More specifically, the degree of non-vanishing ambiguity is increasing in the value of the γ parameter of our prior-selection rule. To illustrate this γ-sensibility of non-vanishing ambiguity, let us first agree on a straightforward definition of the (here: incomplete) "less ambiguous than" relationship in terms of set-inclusion. Given any two sets of probability measures $\mathcal{M}, \mathcal{M}'$ defined on $(\Theta, \mathcal{F})$, we say that $\mathcal{M}$ *expresses (strictly) less ambiguity than* $\mathcal{M}'$ if $\mathcal{M} \subseteq (\subset) \, \mathcal{M}'$.

To focus our analysis, we next impose the following assumption on the *a priori* decision situation in the Ellsberg one-urn experiment.[24]

Assumption 1. Suppose that the set of priors is given as the set of all probability measures on $(\Theta, \mathcal{F})$, i.e., $\mathcal{M}_0 = \triangle^{59}$.

By Assumption 1, we obviously have for any γ that $\Pi_\gamma^\infty \subseteq \mathcal{M}_0$ so that ambiguity must (weakly) decrease through Bayesian learning. The following result shows very concretely that $\gamma' \geq \gamma$ implies $\Pi_\gamma^\infty \subseteq \Pi_{\gamma'}^\infty$ whereby the set-inclusion becomes strict whenever γ' is sufficiently greater than γ. That is, Π_γ^∞ expresses strictly less ambiguity than $\Pi_{\gamma'}^\infty$ if $\gamma' >> \gamma$.

Proposition 4. For any $\theta \in \Theta$,

$$\delta_\theta \in \Pi_\gamma^\infty \tag{102}$$

[24] This assumption corresponds to Gilboa and Schmeidler's (1989, p. 142) 'extreme case'.

if, and only if,

$$\gamma \geq \frac{\frac{1}{3} \cdot \ln \frac{\theta}{90} + \frac{1}{3} \cdot \ln \left(\frac{60-\theta}{90}\right)}{\frac{2}{3} \cdot \ln \frac{1}{3}}. \tag{103}$$

Obviously, the right side of the inequality (103) must take its minimum at the true value $\theta = \theta^* = 30$ so that $\delta_{\theta^*} \in \Pi_\gamma^\infty$ holds for all possible values of γ, including the maximum expected loglikelihood rule where $\gamma = 1$. If the distance $|\theta - \theta^*|$ from the true value increases, however, greater values of γ are required to ensure $\delta_\theta \in \Pi_\gamma^\infty$ whereby these values are, by the symmetry of (103), identical for parameters θ and $\theta' = 60 - \theta$.

To be concrete, Table 1 lists, for all parameter values in Θ, the (approximate) values of γ such that (103) holds with equality.

$\theta; \theta'$	γ
1; 59	2.24014
2; 58	1.93245
3; 57	1.75583
4; 56	1.63296
5; 55	1.5396
6; 54	1.46497
7; 53	1.40332
8; 52	1.35122
9; 51	1.30645
10; 50	1.26751
11; 49	1.23333
12; 48	1.20311
13; 47	1.17627
14; 46	1.15233
15; 45	1.13093

$\theta; \theta'$	γ
16; 44	1.11178
17; 43	1.09466
18; 42	1.07935
19; 41	1.06571
20; 40	1.05361
21; 39	1.04292
22; 38	1.03357
23; 37	1.02548
24; 36	1.01858
25; 35	1.01282
26; 34	1.00816
27; 33	1.00457
28; 32	1.00203
29; 31	1.00051
30	1

Table 1: γ-sensibility of the set of emerging posteriors

The interpretation of Table 1 is straightforward whereby we restrict, for convenience, attention to the subset of emerging posteriors, denoted $\Pi_\gamma^\infty \cap \mathcal{D}$, that only contains emerging Dirac measures.[25] Iff

$$1 \leq \gamma < 1.00051, \tag{104}$$

then

$$\Pi_{\gamma}^{\infty} = \Pi_{\gamma}^{\infty} \cap \mathcal{D} = \{\delta_{\theta^*}\}; \tag{105}$$

iff

$$1.00051 \leq \gamma < 1.00203, \tag{106}$$

then

$$\Pi_{\gamma}^{\infty} \cap \mathcal{D} = \{\delta_{29}, \delta_{\theta^*}, \delta_{31}\}; \tag{107}$$

and so forth until $\gamma \geq 2.24014$ results in

$$\Pi_{\gamma}^{\infty} \cap \mathcal{D} = \{\delta_1, \delta_2, ..., \delta_{59}\}. \tag{108}$$

In this latter case, $\Pi_{\gamma}^{\infty} \cap \mathcal{D} \circ \Phi = \Phi$ so that the decision maker regards all φ_θ in Φ as possible despite the fact that he had observed an unlimited amount of drawings with replacement from the urn.

Note that if, and only if, $\gamma < 1.00051$, ambiguity vanishes in the *a posteriori* decision situation to the effect that the decision maker learns the true value. Furthermore, the decision maker will commit an Ellsberg paradox in the *a posteriori* decision situation if, and only if, $\gamma \geq 1.00051$ because

$$\delta_{29}, \delta_{31} \in \Pi_{\gamma}^{\infty} \tag{109}$$

ensures, by Lemma 1, that the STP violations do not vanish.

6 Related models and an outlook on future research

While the multiple priors decision maker of our single-likelihood environment expresses ambiguity about the index spaces, he is certain that the data is generated by some i.i.d.

[25] Note that, for all θ and $\theta' = 60 - \theta$,

$$D_{KL}(\varphi_{\theta^*}||\varphi_\theta) = D_{KL}(\varphi_{\theta^*}||\varphi_{\theta'})$$

so that there are priors in $\mathcal{M}_0 = \triangle^{59}$, e.g.,

$$0.5\delta_\theta + 0.5\delta_{\theta'},$$

with two different KL-divergence minimizers in their support. By Berk's Theorem, posteriors formed from these priors will not become Dirac measures.

process. To be more specific, recall that Bayesian learning of a Savage (1954) decision maker can be described by the joint index and sample space

$$\left(\Theta \times \Omega^{\infty}, \mathcal{F} \otimes \Sigma^{\infty}, P\right) \tag{110}$$

such that the unique subjective additive probability measure P is pinned down by a unique prior μ_0 as follows: for all $\Theta' \in \mathcal{F}$ and all t,

$$\begin{aligned} P\left(\Theta' \times \left(X_1, ..., X_t\right)\right) &= \sum_{\theta \in \Theta'} \prod_{i=1}^{t} \frac{d\varphi_{\theta}}{dm}\left(X_i\right) \cdot \mu_0\left(\theta\right) & (111) \\ &\equiv P^{\mu_0}\left(\cdot\right) & (112) \end{aligned}$$

As a specific multiple priors generalization of this Savage decision maker, we have considered in this paper a multiple priors decision maker, whom we model via the joint index and sample space

$$\left(\Theta \times \Omega^{\infty}, \mathcal{F} \otimes \Sigma^{\infty}, \Pi^{iid}\right) \tag{113}$$

whereby the additive probability measures in Π^{iid} are pinned down by the multiple priors in $\mathcal{M}_0$ such that

$$\Pi^{iid} = \bigcup_{\mu_0 \in \mathcal{M}_0} \left\{P^{\mu_0}\right\} \tag{114}$$

with P^{μ_0} given by (111).

Instead of the specific i.i.d. multiple priors space (113), one might model Bayesian learning with multiple priors for general spaces

$$\left(\Theta \times \Omega^{\infty}, \mathcal{F} \otimes \Sigma^{\infty}, \Pi\right) \tag{115}$$

where Π stands for an arbitrary set of multiple additive probability measures. For example, the multiple likelihoods environment considered by Epstein and Schneider (2007) weakens the assumption that the decision maker perceives the data as identically distributed whereas it keeps the independence assumption (i.e., the urns might be independently swapped whereby the decision maker cannot observe this swapping). In this ES-2007 multiple-likelihoods environment, the set Π in (115) consists of additive probability measures such that, for all priors $\mu_0 \in \mathcal{M}_0$,

$$P\left(\Theta' \times \left(X_1, X_2, ...\right)\right) = \sum_{\theta \in \Theta'} \prod_{i=1}^{t} \frac{d\varphi_{\theta}^{i}}{dm}\left(X_i\right) \cdot \mu_0\left(\theta\right) \tag{116}$$

with

$$\varphi_{\theta}^{i} \in \Phi\left(\cdot \mid \theta\right) \equiv \left\{\varphi\left(\cdot \mid \theta\right) \mid \varphi \in \Phi\right\}. \tag{117}$$

ES-2007 call $\Phi(\cdot \mid \theta)$ the *multiple likelihoods* set for a given index θ. The question arises in how far our γ-expected maximum loglikelihood rule would also select among different likelihoods and not only among different priors. We therefore regard it an interesting avenue for future research to extend our learning model to a multiple likelihoods environment.

In a different strand of the literature on Bayesian learning under ambiguity, the decision maker is described as a Choquet decision maker whose ambiguity with respect to the joint index and data space is modeled through a non-additive probability measure (e.g., Zimper and Ludwig 2009; Zimper 2011, 2013, Ludwig and Zimper 2014, Groneck, Ludwig, and Zimper 2015). Denote this non-additive probability measure by ν and recall that there exist different perceivable Bayesian update rules according to which a Choquet decision maker may form a conditional non-additive measure $\nu(\cdot \mid \cdot)$ from ν; (cf., e.g., Gilboa and Schmeidler 1993; Sarin and Wakker 1998; Eichberger, Grant, and Kelsey 2006). Such a Choquet Bayesian learner would be modeled for the joint index and sample space

$$(\Theta \times \Omega^{\infty}, \mathcal{F} \otimes \Sigma^{\infty}, \nu(\cdot \mid \cdot)) \tag{118}$$

such that the conditional measure $\nu(\cdot \mid \cdot)$ must specify the Bayesian update rule (resp. rules) that has (resp. have) been applied. More specifically, the Choquet learning models by Zimper and coauthors consider *neo-additive* probability measures (Chateauneuf, Eichberger, and Grant 2007) that are updated either via the *pessimistic*, *optimistic* or *generalized* Bayesian update rule in the light of an objective i.i.d. data process. Such a Choquet decision maker is not only ambiguous about the indices but also about the whole data-generation process. Interestingly, such ambiguity with respect to the data process might result in neo-additive posteriors $\nu(\Theta' \mid X_1, X_2, ...)$ that reflect an increase rather than a decrease in ambiguity whenever the decision maker observes more and more φ_{θ^*}-i.i.d. generated data. Such a possible increase in ambiguity through Bayesian learning is in contrast to this paper's learning model but also to the multiple likelihoods learning model of ES-2007.

It is well-known in the literature that Choquet decision making with respect to conditional neo-additive probability measures can be equivalently described as α-maxmin multiple priors decision making in the sense of Ghirardato, Maccheroni, and Marinacci (2004) if the corresponding multiple priors Bayesian update rules are employed to form sets of multiple posteriors. In future research, we would like to recast neo-additive Choquet Bayesian learning models (118) as Bayesian learning models with multiple priors (115) with the aim to investigate the exact mathematical relationship between these different model classes of Bayesian learning under ambiguity. In particular, the formal relationship between the multiple likelihoods approach of ES-2007, on the one hand, and Choquet Bayesian learning models, on the other hand, is not well-understood yet.

7 Concluding remarks

Nicholls et al.'s (2015) report an experiment in which the number of STP violations does not decline through an increase in statistical information. Motivated by this experimental finding, we have developed a Bayesian learning model with multiple priors such that STP violations do not necessarily vanish. Our approach thereby follows Epstein and Schneider (2007), who convincingly argue that a multiple priors decision maker should test the plausibility of his priors against the observed data. In contrast to the ES-2007 model, however, we consider a more cautious prior selection rule which is governed by a "stubbornness" factor measuring the decision maker's reluctance to revisit his priors. As a potentially interesting feature for future economic applications, the Bayesian learner of our model will end up with the more non-vanishing ambiguity, the more stubborn he is.

Appendix: Formal proofs

Proof of Proposition 1. Step 1. By Definition 5, the expected loglikelihood maximizing prior(s) will be in $\mathcal{M}_{0,\gamma}^{t}$ for all t, i.e.,

$$\arg\max_{\mu_0 \in \mathcal{M}_0} \sum_{\theta\in\Theta} \ln \prod_{i=1}^{t} \frac{d\varphi_\theta}{dm} \cdot \mu_0\left(\theta\right) \subseteq \mathcal{M}_{0,\gamma}^{t}. \tag{119}$$

By a similar formal argument as under Step 3 below, it can be shown that $\mathcal{M}_{0,\gamma}^{\infty}$ (a.s. P_{θ^*}) is never empty since

$$\arg\min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_\theta) \cdot \mu_0\left(\theta\right) \subseteq \mathcal{M}_{0,\gamma}^{\infty} \text{ a.s. } P_{\theta^*}, \tag{120}$$

i.e., the expected Kullback-Leibler divergence minimizers belong asymptotically to the expected γ-log-likelihood maximizers for any value of γ.

Step 2. Observe that any

$$\mu_0' \in \arg\min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_\theta) \cdot \mu_0\left(\theta\right) \tag{121}$$

belongs to (57) if (57) is non-empty.

Step 3. Suppose now that

$$\mu_0' \in \mathcal{M}_{0,\gamma}^{\infty} \tag{122}$$

but

$$\mu_0' \notin \arg\min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_\theta) \cdot \mu_0\left(\theta\right). \tag{123}$$

This is only possible if there exists some subsequence $\{t_k\}_{k\in\mathbb{N}} \subseteq \{t\}_{t\in\mathbb{N}}$ such that

$$\sum_{\theta\in\Theta}\sum_{i=1}^{t_k} \ln\frac{d\varphi_\theta}{dm}\cdot\mu_0'\left(\theta\right) \;\geq\; \gamma\cdot\max_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0\left(\theta\right) \Leftrightarrow \tag{124}$$

$$\sum_{\theta\in\Theta}\frac{1}{t_k}\sum_{i=1}^{t_k} \ln\frac{d\varphi_\theta}{dm}\cdot\mu_0'\left(\theta\right) \;\geq\; \gamma\cdot\max_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0\left(\theta\right) \Rightarrow \tag{125}$$

$$\lim_{t_k\to\infty}\sum_{\theta\in\Theta}\frac{1}{t_k}\sum_{i=1}^{t_k} \ln\frac{d\varphi_\theta}{dm}\cdot\mu_0'\left(\theta\right) \;\geq\; \lim_{t_k\to\infty}\gamma\cdot\max_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0\left(\theta\right). \tag{126}$$

Focus on the l.h.s. term of (126). Because Θ is finite, we can switch the sum and the limit to obtain

$$\lim_{t_k\to\infty}\sum_{\theta\in\Theta}\frac{1}{t_k}\sum_{i=1}^{t_k} \ln\frac{d\varphi_\theta}{dm}\cdot\mu_0'\left(\theta\right) = \sum_{\theta\in\Theta}\lim_{t_k\to\infty}\frac{1}{t_k}\sum_{i=1}^{t_k} \ln\frac{d\varphi_\theta}{dm}\cdot\mu_0'\left(\theta\right). \tag{127}$$

Turn now to the r.h.s. term of (126). We are going to argue, via Berge's (1997) maximum theorem, that we can switch the max and the limit. To this purpose, define the following inner product

$$f\left(y_{t_k},\mu_0\right) \equiv \sum_{\theta\in\Theta}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0\left(\theta\right) \tag{128}$$

$$= y_{t_k}\cdot\mu_0 \tag{129}$$

where

$$y_{t_k}=\left(\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_{\theta_1}}{dm},...,\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_{\theta_n}}{dm}\right) \tag{130}$$

and

$$\mu_0=\left(\mu_0\left(\theta_1\right),...,\mu_0\left(\theta_n\right)\right). \tag{131}$$

Next define the value function of (128) as

$$M\left(y_{t_k}\right)=\max_{\mu_0\in\mathcal{M}_0} f\left(y_{t_k},\mu_0\right). \tag{132}$$

Since $\mathcal{M}_0$ is, as a closed subset of $\triangle^n$, compact and f is continuous, we know from Berge's (1997, p. 116)[26] maximum theorem that the value function $M\left(y_{t_k}\right)$ is continuous. Consequently, if $\lim_{t_k\to\infty} y_{t_k}$ exists, then

$$\lim_{t_k\to\infty}\gamma\cdot M\left(y_{t_k}\right)=\gamma\cdot M\left(\lim_{t_k\to\infty} y_{t_k}\right). \tag{133}$$

In other words, if

$$\lim_{t_k\to\infty}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm} \tag{134}$$

exists (a.s. P_{θ^*}) for all θ, which we will show in a moment, then (a.s. P_{θ^*})

$$\lim_{t_k\to\infty}\gamma\cdot\max_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0\left(\theta\right) \tag{135}$$

$$=\gamma\cdot\max_{\mu_0\in\mathcal{M}_0}\lim_{t_k\to\infty}\sum_{\theta\in\Theta}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0\left(\theta\right) \tag{136}$$

$$=\gamma\cdot\max_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta}\lim_{t_k\to\infty}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0\left(\theta\right). \tag{137}$$

Recall that the law of large numbers implies for the i.i.d.

$$\ln\frac{d\varphi_\theta}{dm}\left(X_1\right),...,\ln\frac{d\varphi_\theta}{dm}\left(X_n\right) \tag{138}$$

[26] Also see p. 570 in Aliprantis and Border (2006).

that

$$\lim_{t_k \to \infty} \frac{1}{t_k} \sum_{i=1}^{t_k} \ln \frac{d\varphi_\theta}{dm} = E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_\theta}{dm} \right] \text{ a.s. } P_{\theta^*} \tag{139}$$

for any θ. By (139) and using (127) and (137), we obtain that (126) is (a.s. P_{θ^*}) equivalent to

$$\sum_{\theta\in\Theta} E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_\theta}{dm} \right] \cdot \mu_0' (\theta) \geq \gamma \cdot \max_{\mu_0 \in \mathcal{M}_0} \sum_{\theta\in\Theta} E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_\theta}{dm} \right] \cdot \mu_0 (\theta) \Leftrightarrow \tag{140}$$

$$\sum_{\theta\in\Theta} E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_\theta}{dm} \right] \cdot \mu_0' (\theta) \geq \gamma \cdot \left(- \min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta\in\Theta} -E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_\theta}{dm} \right] \cdot \mu_0 (\theta) \right) \tag{141}$$

$$\Leftrightarrow$$

$$-\sum_{\theta\in\Theta} E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_\theta}{dm} \right] \cdot \mu_0' (\theta) + \gamma \cdot E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_{\theta^*}}{dm} \right] \tag{142}$$

$$\leq \gamma \cdot \min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta\in\Theta} -E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_\theta}{dm} \right] \cdot \mu_0 (\theta) + \gamma \cdot E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_{\theta^*}}{dm} \right]$$

$$\Leftrightarrow$$

$$\gamma \cdot \sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) d\mu_0' (\theta) - (1-\gamma) \sum_{\theta\in\Theta} E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_\theta}{dm} \right] \cdot \mu_0' (\theta) \tag{143}$$

$$\leq \gamma \cdot \min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) d\mu_0 (\theta)$$

$$\Leftrightarrow$$

$$\sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) \cdot \mu_0' (\theta) \tag{144}$$

$$\leq \min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) d\mu_0 (\theta) + \frac{(1-\gamma)}{\gamma} \sum_{\theta\in\Theta} E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_\theta}{dm} \right] \cdot \mu_0' (\theta) .$$

This proves that

$$\mu_0' \notin \arg \min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) \cdot \mu_0 (\theta) \tag{145}$$

is in $\mathcal{M}_{0,\gamma}^\infty$ if, and only if, μ_0' is in (57).

Step 4. Combining the last argument with Step 1 shows that

$$\arg \min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) \cdot \mu_0 (\theta) = \mathcal{M}_{0,\gamma}^\infty \text{ a.s. } P_{\theta^*} \tag{146}$$

whenever (57) is empty.

Collecting results proves the proposition.□□

Proof of Proposition 2. The only-if part is trivial. By Berk's Theorem 0, we cannot have $\delta_{\hat{\theta}} \in \Pi_{\gamma}^{\infty}$ if $\hat{\theta}$ is not the Kullback-Leibler divergence minimizer for any prior in $\mathcal{M}_0$.

Let us prove the if-part. As a sufficient condition for $\mu_0' \in \mathcal{M}_{0,\gamma}^{\infty}$, we have that

$$\sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_\theta) \cdot \mu_0'(\theta) \leq \frac{(1-\gamma)}{\gamma} \sum_{\theta\in\Theta} E_{\varphi_{\theta^*}}\left[\ln \frac{d\varphi_\theta}{dm}\right] \cdot \mu_0'(\theta). \tag{147}$$

Next observe that

$$\lim_{\gamma\to\infty} \frac{(1-\gamma)}{\gamma} \sum_{\theta\in\Theta} E_{\varphi_{\theta^*}}\left[\ln \frac{d\varphi_\theta}{dm}\right] \cdot \mu_0'(\theta) = -\sum_{\theta\in\Theta} E_{\varphi_{\theta^*}}\left[\ln \frac{d\varphi_\theta}{dm}\right] \cdot \mu_0'(\theta) \tag{148}$$

as well as

$$\begin{aligned}\sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_\theta) \cdot \mu_0'(\theta) &= \sum_{\theta\in\Theta} E_{\varphi_{\theta^*}}\left[\ln \frac{d\varphi_{\theta^*}}{dm}\right] \cdot \mu_0'(\theta) - \sum_{\theta\in\Theta} E_{\varphi_{\theta^*}}\left[\ln \frac{d\varphi_\theta}{dm}\right] \cdot \mu_0'(\theta) \\ &< -\sum_{\theta\in\Theta} E_{\varphi_{\theta^*}}\left[\ln \frac{d\varphi_\theta}{dm}\right] \cdot \mu_0'(\theta) \end{aligned} \tag{149}$$

since, by (71),

$$\sum_{\theta\in\Theta} E_{\varphi_{\theta^*}}\left[\ln \frac{d\varphi_{\theta^*}}{dm}\right] \cdot \mu_0'(\theta) < 0. \tag{150}$$

Consequently, we can always find γ large enough such that

$$\sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_\theta) \cdot \mu_0'(\theta) < \frac{(1-\gamma)}{\gamma} \sum_{\theta\in\Theta} E_{\varphi_{\theta^*}}\left[\ln \frac{d\varphi_\theta}{dm}\right] \cdot \mu_0'(\theta), \tag{151}$$

which proves that $\mu_0' \in \mathcal{M}_{0,\gamma}^{\infty}$.

In other words, any $\mu_0' \in \mathcal{M}_0$ will belong to $\mathcal{M}_{0,\gamma}^{\infty}$ if γ is chosen is sufficiently large. If such μ_0' has $\hat{\theta}$ as the unique Kullback-Leibler divergence minimizer in its support, we obtain (72) by Berk's Theorem 0.□□

Proof of Proposition 3. Step 1. Consider the *a priori* decision situation. Analogous to the derivation of Lemma 1, we have that

$$MEU\left(f_E h, \mathcal{M}_0 \circ \Phi\right) = \frac{1}{3} \text{ and } MEU\left(g_E h', \mathcal{M}_0 \circ \Phi\right) = \frac{2}{3}. \tag{152}$$

Further, note that

$$\begin{aligned} MEU\left(g_E h, \mathcal{M}_0 \circ \Phi\right) &\leq MEU\left(g_E h, \sum_{\theta\in\Theta} \varphi_\theta\left(\omega\right) \mu_0'\left(\theta\right)\right) && (153)\\ &\leq EU\left(g_E h, \sum_{\theta\in\Theta} \varphi_\theta\left(\omega\right) \delta_{29}\right) && (154)\\ &= EU\left(g_E h, \varphi_{29}\right) && (155)\\ &= \frac{29}{90} < \frac{1}{3} && (156) \end{aligned}$$

as well as

$$\begin{aligned} MEU\left(f_E h', \mathcal{M}_0 \circ \Phi\right) &\leq MEU\left(f_E h', \sum_{\theta\in\Theta} \varphi_\theta\left(\omega\right) \mu_0''\left(\theta\right)\right) && (157)\\ &\leq EU\left(f_E h', \sum_{\theta\in\Theta} \varphi_\theta\left(\omega\right) \delta_{31}\right) && (158)\\ &= EU\left(f_E h', \varphi_{31}\right) && (159)\\ &= \frac{1}{3} + \frac{29}{90} < \frac{2}{3}. && (160) \end{aligned}$$

Consequently, the inequalities (98)-(99) hold.

Step 2. Consider the *a posteriori* decision situation. Note that

$$MEU\left(f_E h, \Pi_\gamma^\infty \circ \Phi\right) = \frac{1}{3} \text{ and } MEU\left(g_E h', \Pi_\gamma^\infty \circ \Phi\right) = \frac{2}{3}. \quad (161)$$

The specifications of μ_0' and μ_0'' imply, by Proposition 2, for some sufficiently large γ the existence of some

$$\delta_{\theta'}, \delta_{\theta''} \in \Pi_\gamma^\infty \quad (162)$$

such that $\theta' < 30 < \theta''$. Consequently,

$$\begin{aligned} MEU\left(g_E h, \Pi_Z^\infty \circ \Phi\right) &\leq EU\left(g_E h, \sum_{\theta\in\Theta} \varphi_\theta\left(\omega\right) \delta_{\theta'}\right) && (163)\\ &\leq EU\left(g_E h, \sum_{\theta\in\Theta} \varphi_\theta\left(\omega\right) \delta_{29}\right) && (164)\\ &< \frac{1}{3} && (165) \end{aligned}$$

as well as

$$MEU\left(f_E h', \Pi_Z^\infty \circ \Phi\right) \leq EU\left(f_E h', \sum_{\theta\in\Theta} \varphi_\theta(\omega)\,\delta_{\theta''}\right) \tag{166}$$

$$\leq EU\left(f_E h', \sum_{\theta\in\Theta} \varphi_\theta(\omega)\,\delta_{31}\right) \tag{167}$$

$$< \frac{2}{3}, \tag{168}$$

which proves the inequalities (100)-(101).□□

Proof of Proposition 4. By the proof of Proposition 1 (cf., inequality (140) as well as Step 2.), $\mu_0' \in \mathcal{M}_{0,\gamma}^\infty$ if, and only if,

$$\sum_{\theta'\in\Theta} E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_{\theta'}}{dm}\right]\cdot\mu_0'(\theta') \geq \gamma\cdot\max_{\mu_0\in\mathcal{M}_0}\sum_{\theta'\in\Theta} E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_{\theta'}}{dm}\right]\cdot\mu_0(\theta'). \tag{169}$$

Since, by assumption, $\mathcal{M}_0 = \triangle^{59}$, we have, for any $\theta\in\Theta$,

$$\delta_\theta \in \Pi_\gamma^\infty \tag{170}$$

if, and only if,

$$E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_{\theta'}}{dm}\right]\cdot\delta_\theta \geq \gamma\cdot E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_{\theta'}}{dm}\right]\cdot\delta_{\theta^*}$$

$$\Leftrightarrow$$

$$E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_{\theta}}{dm}\right] \geq \gamma\cdot E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_{\theta^*}}{dm}\right] \tag{171}$$

$$\Leftrightarrow$$

$$\frac{E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_{\theta}}{dm}\right]}{E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_{\theta^*}}{dm}\right]} \leq \gamma \tag{172}$$

$$\Leftrightarrow$$

$$\frac{d\varphi_{\theta^*}(\omega_2)\cdot\ln d\varphi_\theta(\omega_2) + \left(\frac{2}{3} - d\varphi_{\theta^*}(\omega_2)\right)\cdot\ln\left(\frac{2}{3} - d\varphi_\theta(\omega_2)\right)}{d\varphi_{\theta^*}(\omega_2)\cdot\ln d\varphi_{\theta^*}(\omega_2) + \left(\frac{2}{3} - d\varphi_{\theta^*}(\omega_2)\right)\cdot\ln\left(\frac{2}{3} - d\varphi_{\theta^*}(\omega_2)\right)} \leq \gamma \tag{173}$$

$$\Leftrightarrow$$

$$\frac{\frac{1}{3}\cdot\ln\frac{\theta}{90} + \frac{1}{3}\cdot\ln\left(\frac{60-\theta}{90}\right)}{\frac{2}{3}\cdot\ln\frac{1}{3}} \leq \gamma, \tag{174}$$

which proves the proposition.□□

References

Aliprantis, D. C., and K. Border (2006) *Infinite Dimensional Analysis* (2nd edn.) Springer: Berlin.

Anscombe, F. J., and R. J. Aumann (1963) "A Definition of Subjective Probability," *Annals of American Statistics* **34**, 199-205.

Berge, C. (1997) *Topological Spaces*, Dover Publications, Inc.: New York. (An "unabridged and unaltered republication" of the original English publication from 1963.)

Berk, R. H. (1966) "Limiting Behaviour of Posterior Distributions when the Model is Incorrect," *The Annals of Mathematical Statistics* **31**, 51–58.

Billingsley, P. (1995) *Probability and Measure*, John Wiley and & Sons, Inc.: New York.

Chateauneuf, A., Eichberger, J., and S. Grant (2007) "Choice under Uncertainty with the Best and Worst in Mind: Neo-additive Capacities," *Journal of Economic Theory* **127**, 538-567.

Chung, K. L., and W. H. J. Fuchs (1951) "On the Distribution of Values of Sums of Random Variables," *Memoirs of the American Mathematical Society* **6**, 1-12. (Reprinted in: F. AitSaliha, E. Hsu, and R. Williams [eds.], *Selected Works of Kai Lai Chung*, 2008, pp.157-168, World Scientific: New Jersey.)

Diaconis, P. and D. Freedman (1986) "On the Consistency of Bayes Estimates," *The Annals of Statistics* **14**, 1-26.

Dow, J., Madrigal, V., and S. R. Werlang (1990), "Preferences, Common Knowledge, and Speculative Trade," mimeo.

Eichberger, J., Grant, S., and D. Kelsey (2006) "Updating Choquet Beliefs," *Journal of Mathematical Economics* **43**, 888-899.

Ellsberg, D. (1961) "Risk, Ambiguity and the Savage Axioms," *Quarterly Journal of Economics* **75**, 643-669.

Epstein L. G., and M. Schneider (2007) "Learning under Ambiguity," *Review of Economic Studies* **74**, 1275–1303.

Doob, J. L. (1949) "Application of the Theory of Martingales," In: *Le Calcul des Probabilites et ses Applications*, Colloques Internationaux du Centre National de la Recherche Scientifique 13, pp. 23–27. Paris: CNRS.

Ghirardato, P., Maccheroni, F., and M. Marinacci (2004) “Differentiating Ambiguity and Ambiguity Attitude,” *Journal of Economic Theory* **118**, 133–173.

Gilboa, I. (1987) “Expected Utility with Purely Subjective Non-Additive Probabilities,” *Journal of Mathematical Economics* **16**, 65-88.

Gilboa, I., and D. Schmeidler (1989) “Maxmin Expected Utility with Non-Unique Priors,” *Journal of Mathematical Economics* **18**, 141-153.

Gilboa, I., and D. Schmeidler (1993) “Updating Ambiguous Beliefs,” *Journal of Economic Theory* **59**, 33-49.

Gosh, J. K., and R. V. Ramamoorthi (2003) *Bayesian Nonparametrics*, Springer Verlag Inc.: New York.

Groneck, M., Ludwig, A., and A. Zimper (2015) “A Life-cycle Model with Ambiguous Survival Beliefs,” Working paper.

Halevy, Y. (2004) “The Possibility of Speculative Trade between Dynamically Consistent Agents”, *Games and Economic Behavior* **46**, 189-198.

Harrison, M., and D. Kreps (1978) “Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations,” *Quarterly Journal of Economics* **42**, 323–36.

Jaffray, J-Y. (1994) “Dynamic Decision Making with Belief Functions,” In: R. R. Yager, M. Fedrizzi, and J. Kacprzyk [eds.], *Advances in the Dempster-Shafer Theory of Evidence*, pp.331–352. Wiley, New York.

Kleijn, B. J. K., and A. W. van der Vaart (2006) “Misspecification in Infinite-dimensional Bayesian Statistics,” *The Annals of Statistics* **34**, 837–877.

Kullback S., and R. A. Leibler (1951) “On Information and Sufficiency,” *The Annals of Mathematical Statistics* **22**, 79–86.

Lijoi, A., Pruenster, I. and S.O. Walker (2004) “Extending Doob’s Consistency Theorem to Nonparametric Densities,” *Bernoulli*, **10**, 651-663.

Ludwig, A., and A. Zimper (2014) “Biased Bayesian Learning with an Application to the Risk-free Rate Puzzle,” *Journal of Economic Dynamics and Control* **39**, 79-97.

Marinacci, M. (2002) “Learning from Ambiguous Urns,” *Statistical Papers* **43**, 145-151.

Mehra, R., and E. C. Prescott (1985) “The Equity Premium: A Puzzle,” *Journal of Monetary Economics* **15**, 145-161.

Mehra, R., and E. C. Prescott (2003) "The Equity Premium in Retrospect," in: G.M. Constantinides, M. Harris and R.M. Stulz, eds., *Handbook of the Economics and Finance*. Amsterdam: Elsevier, 808-887.

Muth, J. F. (1961) "Rational Expectations and the Theory of Price Movements," *Econometrica* **29**, 315-335.

Nicholls, N., Romm, A. T., and A. Zimper (2015) "The Impact of Statistical Learning on Violations of the Sure-thing Principle," *Journal of Risk and Uncertainty*, forthcoming.

Sarin, R., and P. P. Wakker (1998) "Revealed Likelihood and Knightian Uncertainty," *Journal of Risk and Uncertainty* **16**, 223-250.

Savage, L. J. (1954) *The Foundations of Statistics*, John Wiley and & Sons, Inc.: New York, London, Sydney.

Schmeidler, D. (1989) "Subjective Probability and Expected Utility without Additivity," *Econometrica* **57**, 571-587.

Wakker, P. P. (2010), *Prospect Theory for Risk and Ambiguity*, Cambridge University Press.

Werner, J. (2014) "Speculative Trade under Ambiguity," Working paper.

Wu, G., and R. Gonzalez (1999) "Nonlinear Decison Weighst in Choice under Uncertainty," *Management Science* **45**, 74-85.

Zimper, A. (2009) "Half Empty, Half Full and Why We Can Agree to Disagree Forever," *Journal of Economic Behavior and Organization* **71**, 283-299.

Zimper, A. (2011) "Do Bayesians Learn their Way out of Ambiguity?," *Decision Analysis* **8**, 269-285.

Zimper, A. (2013) "The Emergence of "Fifty-Fifty" Probability Judgments through Bayesian Updating under Ambiguity," *Fuzzy Sets and Systems* **223**, 72-88.

Zimper, A., and A. Ludwig (2009) "On Attitude Polarization under Bayesian Learning with Non-additive Beliefs," *Journal of Risk and Uncertainty* **39**, 181-212.