Bayesian learning with multiple priors and non-vanishing ambiguity*

Alexander Zimper† Wei Ma‡

July 24, 2016

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 survive a prior selection rule which discriminates, with probability one, against priors whose expected Kullback-Leibler divergence from the 'truth' is too far off from the minimal expected Kullback-Leibler divergence over all priors. The "stubbornness" parameter of our prior selection rule thereby governs how much ambiguity will remain 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, Daniele Pennesi, an anonymous referee and, especially, Jan Werner for helpful comments and suggestions. Financial support from ERSA (Economic Research Southern Africa) is gratefully acknowledged.

†Corresponding author. Department of Economics, University of Pretoria, and the Kiel Institute for the World Economy. E-mail: alexander.zimper@up.ac.za

‡International Business School, Suzhou Xi'an Jiaotong-Liverpool University, and Department of Economics, University of Pretoria. E-mail: Wei.Ma@xjtlu.edu.cn

1 Introduction

1.1 Motivation

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)

This paper constructs a model of Bayesian learning with multiple priors that can accommodate this empirical finding in a plausible way. As the key feature of our model, the decision maker rejects priors in the light of observed data by an application of the

[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.

γ*-maximum expected loglikelihood* prior selection rule, which we newly introduce to the literature. We interpret γ as a "stubbornness" parameter because it measures the decision maker's reluctance to dismiss prior beliefs in the light of new data. While ambiguity typically decreases through statistical learning in our model, it does not necessarily vanish to the effect that STP violations may occur after arbitrarily many statistical observations.

1.2 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. A standard Bayesian learner would resolve 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 indices in the prior's support that minimize the *Kullback-Leibler (1951) divergence* 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,

[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), and Lijoi, Pruenster and Walker (2004).

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.

1.3 New results

Our approach follows ES-2007 in that we also assume that a multiple priors decision maker applies a prior selection rule according to which he might reject priors as implausible in the light of the observed data. We argue, however, that the ES-2007 α-expected maximum likelihood rule might be too strong for some decision makers because it implies vanishing ambiguity in the single-urn environment. 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, our 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.

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

1. In the limit, posteriors are only formed from priors whose expected Kullback-Leibler divergence from the true index/measure is γ-close to the minimal expected Kullback-Leibler divergence from the true index/measure over all priors.

2. For the special case of the maximum expected loglikelihood rule (i.e., $\gamma = 1$), posteriors are only formed from priors that minimize the expected Kullback-Leibler divergence from the true index/measure. In contrast to ES-2007's α-expected maximum likelihood rule, these posteriors are not necessarily concentrating at the true index/measure even if there is a well-specified prior.

3. Greater values of the stubbornness parameter γ imply greater sets of posteriors in the limit 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.

4. In line with the experimental findings of Nicholls et al. (2015), we can therefore generate non-vanishing STP violations for sufficiently large values of the stubbornness parameter γ.

The remainder of the paper is organized as follows. Motivated by the experimental findings of Nicholls et al. (2015), we present in Section 2 a wish list for a prior selection rule. Section 3 recalls the formal set-up of Bayesian learning with a unique prior. Section 4 extends this framework to Bayesian learning with multiple priors. In Section 5 we derive new results for Bayesian learning with multiple priors under the γ-maximum expected loglikelihood prior selection rule. Section 6 revisits the Ellsberg one-urn experiment whereby we illustrate the possibility of non-vanishing STP violations. Section 7 concludes with an outlook on possible economic applications.

2 The Ellsberg experiment and a wish list for a prior selection rule

As our theoretical motivation, we want to explore a plausible alternative to the existing Bayesian learning models with multiple priors proposed by M-2002 and by ES-2007, respectively. As our empirical motivation, we would like to accommodate the experimental findings of Nicholls et al. (2015) in a convincing way. This section elaborates in some detail on this empirical motivation.

2.1 STP violations in the static Ellsberg one-urn experiment

As our point of departure, let us recall the original Ellsberg one-urn experiment as well as the proposal of Gilboa and Schmeidler (1989) to explain STP violations through *max min expected utility* (MEU).

Consider some state space Ω and some σ-algebra Σ. Savage (1954) describes a decision maker who has preferences $\succeq$ over Savage acts which are Σ-measurable mappings from the 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,

$$f \succeq g \Leftrightarrow \int_{\omega \in \Omega} u\left(f\left(\omega\right)\right) d\varphi \geq \int_{\omega \in \Omega} u\left(g\left(\omega\right)\right) d\varphi, \tag{1}$$

whereby the subjective probability measure φ defined on (Ω, Σ) and the utility function $u : Z \to \mathbb{R}$ are uniquely[3] 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

$$EU(f, \varphi) \equiv \int_{\omega \in \Omega} u(f(\omega))\, d\varphi. \tag{2}$$

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

$$f_E h \succeq g_E h \;\Leftrightarrow\; f_E h' \succeq g_E h' \tag{3}$$

whereby

$$f_E h(\omega) = \begin{cases} f(\omega) & \text{for } \omega \in E \\ h(\omega) & \text{for } \omega \in \neg E \end{cases} \tag{4}$$

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

$$\Omega = \{\omega_1, \omega_2, \omega_3\} \tag{5}$$

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 $E \equiv \{\omega_1, \omega_2\}$

	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

$$f_E h \succ g_E h \text{ and } g_E h' \succ f_E h'. \tag{6}$$

Note that this 'Ellsberg paradox' (6) constitutes a violation of the STP (3) and can therefore not be accommodated by SEU theory.

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

To accommodate the Ellsberg paradox (6), Gilboa and Schmeidler (1989) propose a *maxmin expected utility with non-unique prior* (=MEU) representation such that, for all Savage acts f, g,

$$f \succeq g \Leftrightarrow \min_{\varphi \in \mathcal{P}} \int_{\omega \in \Omega} u(f(\omega))\, d\varphi \geq \min_{\varphi \in \mathcal{P}} \int_{\omega \in \Omega} u(g(\omega))\, d\varphi \tag{7}$$

for some non-empty set of probability measures $\mathcal{P}$.[4] 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. 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 rather extreme[5], we follow here Gilboa and Schmeidler (1989) and consider a MEU decision maker. In what follows, we write

$$MEU(f, \mathcal{P}) \equiv \min_{\varphi \in \mathcal{P}} EU(f, \varphi) \tag{8}$$

for 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 (6) suppose (quite naturally) that $\mathcal{P}$ is non-empty and satisfies

$$\mathcal{P} \subseteq \left\{ \varphi = \left(\frac{1}{3}, \varphi(\omega_2), \frac{2}{3} - \varphi(\omega_2) \right) \mid \varphi(\omega_2) \in \left[0, \frac{2}{3}\right] \right\} \tag{9}$$

That is, $\mathcal{P}$ only contains probability measures φ defined on $\left(\Omega, 2^{\Omega}\right)$ that attach probability $\frac{1}{3}$ to the event that a red ball will be drawn. Without loss of generality, set $u(z) = z$ for $z \in \{0, 1\}$. By (9),

$$MEU(f_E h, \mathcal{P}) = \frac{1}{3} \text{ and } MEU(g_E h', \mathcal{P}) = \frac{2}{3}. \tag{10}$$

Note that

$$MEU(g_E h, \mathcal{P}) < \frac{1}{3} \tag{11}$$

[4]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 $\mathcal{P}$.

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

if, and only if, there is some $\varphi \in \mathcal{P}$ such that $\varphi(\omega_2) < \frac{1}{3}$. Conversely, we have that

$$MEU(f_E h', \mathcal{P}) < \frac{2}{3} \tag{12}$$

if, and only if, there is some $\varphi' \in \mathcal{P}$ such that $\varphi'(\omega_2) > \frac{1}{3}$.

Collecting the above arguments gives us the following sufficient and necessary condition for STP violations in the Ellsberg one-urn experiment by an MEU decision maker.

Proposition 1. The MEU decision maker violates the STP, i.e.,

$$MEU(f_E h, \mathcal{P}) > MEU(g_E h, \mathcal{P}), \tag{13}$$

$$MEU(g_E h', \mathcal{P}) > MEU(f_E h', \mathcal{P}), \tag{14}$$

if, and only if, there are $\varphi, \varphi' \in \mathcal{P}$ such that

$$\varphi(\omega_2) < \frac{1}{3} \text{ and } \varphi'(\omega_2) > \frac{1}{3}. \tag{15}$$

2.2 Non-vanishing STP violations in the Ellsberg one-urn experiment with statistical information

Nicholls et al. (2015) report an experimental study in which subjects had to make 30 consecutive choices, structured as 15 choice pairs, such that every choice pair resembled a payoff-variant of Ellsberg's static one-urn experiment. The elicited choices for any given pair were thus either consistent with the STP or they constituted a violation thereof. Whereas the test group received increasing statistical information about the urn's composition in the form of drawings with replacement after each choice, the control group did not receive any such information. At the end of the experiment each subject in the test group had thus observed statistical information in the form of 30 data points independently generated by the true distribution corresponding to the urn's composition.

Since a Bayesian decision maker would use this statistical information to update his prior belief(s) about the composition of the urn, one would intuitively expect that this decision maker's ambiguity should decrease in the amount of the statistical information he received. In line with this intuition, the working hypothesis of Nicholls et al. (2015) was that the occurrence of STP violations should become less frequent for the test group than for the control group over the course of the experiment. Contrary to this hypothesis, however, the experiment did not show a decrease in STP violations for the test group compared to the control group.

Even under the (strong) assumption that the subjects resembled Bayesian learners with multiple priors, there might exist a number of possible 'explanations' for the experimental findings of Nicholls et al. (2015). For example, there might not have been enough statistical information for driving out STP violations for the test group. Nevertheless, a model of Bayesian learning for which ambiguity does not necessarily vanish looks like a better candidate for 'explaining' the experimental finding of non-decreasing STP violations for the whole test group than a model for which STP violations would completely vanish in the long run.

To formally describe non-vanishing STP violations, we now recast the Ellsberg one-urn experiment within a set-up of increasing statistical information about the urn's composition. As before let $\Omega = \{\omega_1, \omega_2, \omega_3\}$ and define the relevant σ-algebra as the powerset of Ω, i.e., $\Sigma = 2^{\Omega}$. Next consider the following set of indexed probability measures defined on (Ω, Σ)

$$\Phi = \left\{ \varphi_\theta = \left(\frac{1}{3}, \varphi_\theta(\omega_2), \frac{2}{3} - \varphi_\theta(\omega_2) \right) \mid \varphi_\theta(\omega_2) = \frac{\theta}{90}, \theta \in \Theta \right\} \tag{16}$$

where Θ denotes an *index set* such that

$$\Theta = \{1, ..., 59\}. \tag{17}$$

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.[6] That is, Φ contains all probability measures with full support on Ω that correspond to all possible ratios of black versus yellow balls in the Ellsberg urn. Denote by $\theta^* \in \Theta$ the actual number of black balls in the urn so that the "true" probability measure φ_{θ^*} determines the distribution according to which balls are drawn from the urn. By assumption, the decision maker does not have any knowledge about the urn's actual composition (i.e., about the true index value) beyond the trivial fact that $\theta^* \in \Theta$.

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. If the decision maker was a standard Savage decision maker he would resolve his *a priori* uncertainty about the urn's composition by a unique *prior* (=additive probability measure), denoted μ_0, defined on the index space $(\Theta, \mathcal{F})$ where $\mathcal{F}$ denotes the powerset of Θ. More generally, we consider a multiple priors decision maker who resolves his uncertainty about the urn's composition by a non-empty set of priors, denoted $\mathcal{M}_0$,

[6] 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 a prior attaches probability one to zero yellow balls in the urn but the decision maker observes a yellow ball drawn from the urn.

defined on the index space $(\Theta, \mathcal{F})$. For a MEU decision maker the set of probability measures $\mathcal{P}$ in (8) 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,

$$MEU\left(f, \mathcal{M}_0 \circ \Phi\right) = \min_{\varphi^0 \in \mathcal{M}_0 \circ \Phi} \sum_{\omega \in \Omega} u\left(f\left(\omega\right)\right) \varphi^0\left(\omega\right) \tag{18}$$

where

$$\varphi^0\left(\omega\right) \equiv \sum_{\theta \in \Theta} \varphi_\theta\left(\omega\right) \mu_0\left(\theta\right). \tag{19}$$

Note that this *a priori* MEU decision maker reduces to an *a priori* SEU decision maker if, and only if, the set of priors $\mathcal{M}_0$ contains a unique prior.

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. After receiving this statistical information, a multiple priors decision maker resolves (with probability one) his uncertainty about the true parameter value through some set of emerging posteriors defined on the index space $(\Theta, \mathcal{F})$. For the moment being, let us denote this set of posteriors by Π^∞ with generic element π^∞. In this *a posteriori* decision situation, the set of probability measures $\mathcal{P}$ in (8) then becomes the set of reduced compound measures $\mathcal{P} = \Pi^\infty \circ \Phi$, defined on (Ω, Σ), implying, for any Savage act f,

$$MEU\left(f, \Pi^\infty \circ \Phi\right) = \min_{\varphi^\infty \in \Pi^\infty \circ \Phi} \sum_{\omega \in \Omega} u\left(f\left(\omega\right)\right) \varphi^\infty\left(\omega\right) \tag{20}$$

where

$$\varphi^\infty\left(\omega\right) \equiv \sum_{\theta \in \Theta} \varphi_\theta\left(\omega\right) \pi^\infty\left(\theta\right). \tag{21}$$

Definition 1. We speak of non-vanishing STP violations if, and only if, the decision maker violates the STP 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,

$$MEU\left(f_E h, \mathcal{M}_0 \circ \Phi\right) > MEU\left(g_E h, \mathcal{M}_0 \circ \Phi\right) \text{ and} \tag{22}$$

$$MEU\left(g_E h', \mathcal{M}_0 \circ \Phi\right) > MEU\left(f_E h', \mathcal{M}_0 \circ \Phi\right) \tag{23}$$

as well as

$$MEU\left(f_E h, \Pi^\infty \circ \Phi\right) > MEU\left(g_E h, \Pi^\infty \circ \Phi\right) \text{ and} \tag{24}$$

$$MEU\left(g_E h', \Pi^\infty \circ \Phi\right) > MEU\left(f_E h', \Pi^\infty \circ \Phi\right). \tag{25}$$

An *a priori* MEU decision maker becomes an *a posteriori* SEU decision maker if, and only if, there is a unique emerging posterior in the set Π^{∞}. Non-vanishing STP violations can thus only be accommodated by a Bayesian learning model that allows for multiple emerging posteriors in Π^{∞} for a given set of multiple priors $\mathcal{M}_0$.

2.3 A wish list for a prior selection rule

To generate non-vanishing STP violations is actually quite trivial for a model of Bayesian learning such that emerging posteriors are formed from **all** priors in $\mathcal{M}_0$. To see this, denote by δ_θ the *Dirac measure* that attaches probability one to the index value $\theta \in \Theta$ and suppose that $\delta_1, \delta_{59} \in \mathcal{M}_0$. That is, one prior attaches probability one to the occurrence of one black ball in the urn whereas another prior attaches probability one to 59 black balls. For these degenerate priors we have that

$$\varphi_1^0(\omega_2) = \sum_{\theta\in\Theta} \varphi_\theta(\omega_2)\,\delta_1 = \frac{1}{90} < \frac{1}{3}, \tag{26}$$

$$\varphi_{59}^0(\omega_2) = \sum_{\theta\in\Theta} \varphi_\theta(\omega_2)\,\delta_{59} = \frac{59}{90} > \frac{1}{3} \tag{27}$$

so that the STP is, by Proposition 1, violated in the *a priori* decision situation. Because Bayesian updating of degenerate priors does not change these priors at all, we end up with emerging posteriors $\delta_1, \delta_{59} \in \Pi^{\infty}$. By the same argument that gave us (26) and (27), we obtain

$$\varphi_1^{\infty}(\omega_2) = \frac{1}{90}, \tag{28}$$

$$\varphi_{59}^{\infty}(\omega_2) = \frac{59}{90}. \tag{29}$$

Consequently, the STP is also violated in the *a posteriori* decision situation, which gives us non-vanishing STP violations.

The above updating of all priors in $\mathcal{M}_0$ is, however, highly problematic because it precludes the possibility that a decision maker might critically revisit his multiple priors in the light of new data. Suppose, for example, that we have as true index value $\theta^* = 30$ and that one of the decision maker's priors attaches probability one to this true value, i.e., $\delta_{30} \in \mathcal{M}_0$. In the *a posteriori* decision situation, a decision maker will have observed with probability one that a black ball was drawn with frequency $\frac{1}{3}$. We regard it as plausible–at least for some decision makers–that the 'correct' prior δ_{30} should be able to drive out the highly incorrect priors δ_1 and δ_{59} as it is highly unlikely that either of these priors has generated the observed data.

For this reason, we follow ES-2007 who argue that Bayesian learning with multiple priors should be a combination of standard Bayesian updating with a *prior selection rule*. The decision maker would then not form posteriors from all priors in $\mathcal{M}_0$ but only from those priors that survive a data likelihood test specified by a prior selection rule. The *α-expected maximum likelihood* prior selection rule proposed by ES-2007 is indeed able to drive out the incorrect priors δ_1 and δ_{59} in favor of the correct prior δ_{30}. As a drawback, however, the ES-2007 prior selection rule is so strong that it can never generate non-vanishing STP violations because it results in a unique emerging posterior (here δ_{30}).

In the light of the above considerations, we state the following wish list for a prior selection rule.

1. As ES-2007, we would like to come up with a plausible prior selection rule that allows the decision maker to reject priors in favor of other priors which are more likely to have generated the observed data.

2. In contrast to ES-2007 (and in line with the experimental findings of Nicholls et al. 2015), however, we want to keep this prior selection rule sufficiently weak to allow for the possibility of non-vanishing STP violations.

3. Preferably, the prior selection rule should also allow for the possibility that different decision makers may exhibit different learning behavior in the limit; that is, the question whether two decision makers who start out with the same set of priors will exhibit non-vanishing STP violations should be determined by some parameter that comes with the prior selection rule.

3 Preliminaries

3.1 Formal set-up

This section develops our general set-up which will nest the Ellsberg one-urn experiment with increasing statistical information as a special case. 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 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 set of probability measures on (Ω, Σ) that are indexed by $\theta \in \Theta$, i.e.,

$$\Phi = \{\varphi_\theta \mid \theta \in \Theta\}. \tag{30}$$

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.[7]

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

$$E\left[I_{X^{-1}(B)}, \varphi_\theta\right] = \int_\Omega I_B(\omega)\, d\varphi_\theta(\omega) \tag{31}$$

$$= \varphi_\theta(B) \tag{32}$$

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 $\varphi_\theta \in \Phi$,

$$\varphi_\theta((a,b]) = \varphi_\theta(X = b) - \varphi_\theta(X = a) \text{ for all } (a,b] \in \Sigma, \tag{33}$$

in the continuous case; and

$$\varphi_\theta(\{\omega_m\}) = \varphi_\theta(X = b) - \varphi_\theta(X = a) \text{ for all } \omega_m \in \Omega \tag{34}$$

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.[8]

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 $B \in \Sigma$,

$$\varphi_\theta(B) = \int_{\omega \in B} \frac{d\varphi_\theta}{dm}(\omega)\, dm. \tag{35}$$

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

[8] 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, $\frac{d\varphi_\theta}{dm} : \mathbb{R} \to \mathbb{R}_+$ stands for the familiar probability density function (=pdf) such that

$$\varphi_\theta((a,b]) = \int_{x\in(a,b]} \frac{d\varphi_\theta}{dm}(x)\,dx. \tag{36}$$

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

$$\begin{aligned}
\varphi_\theta(\omega) &= \int_{\omega'\in\{\omega\}} \frac{d\varphi_\theta}{1}(\omega')\,1 &(37)\\
&= \frac{d\varphi_\theta}{1}(\omega) &(38)\\
&= d\varphi_\theta(\omega). &(39)
\end{aligned}$$

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

$$\varphi_\theta(X=1) = \varphi_\theta(\omega_1) = \theta. \tag{40}$$

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

$$X_t (..., \omega_t, ...) = X(\omega) \text{ for } \omega_t = \omega. \tag{41}$$

$(\Omega^\infty, \Sigma^\infty)$ 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 $(\Theta, \mathcal{F})$.

3.2 Bayesian learning with a unique prior

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. In this subsection we recall the standard case of a Bayesian decision maker who holds a unique prior $\mu_0 \in \triangle^n$ defined on the index space $(\Theta, \mathcal{F})$ where $\triangle^n$ denotes the n-simplex equipped with the Euclidean topology.[9] Through Bayesian updating we obtain the (conditional) probability space $\left(\Theta, \mathcal{F}, \pi^t_{\mu_0}\right)$ such that one version of the posterior $\pi_{\mu_0}(\cdot \mid X_1, ..., X_t) \in \triangle^n$, formed from the prior μ_0 after observing a data sample drawn from $X_1, ..., X_t$, is formally given as

$$\pi_{\mu_0}(\Theta' \mid X_1, ..., X_t) = \frac{\sum_{\theta \in \Theta'} \prod_{i=1}^{t} \frac{d\varphi_\theta}{dm}(X_i) \cdot \mu_0(\theta)}{\sum_{\theta \in \Theta} \prod_{i=1}^{t} \frac{d\varphi_\theta}{dm}(X_i) \cdot \mu_0(\theta)} \tag{42}$$

[9]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 $\Theta' \in \mathcal{F}$,

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

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 $\pi_{\mu_0}(\cdot \mid X_1, ..., X_t)$ converges with probability one to some *emerging*[10] *posterior* $\pi^{\infty}_{\mu_0}$, i.e.,

$$\pi^{\infty}_{\mu_0} \equiv \lim_{t\to\infty} \pi_{\mu_0}(\cdot \mid X_1, ..., X_t) \text{ a.s.} \tag{43}$$

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[11], the emerging posterior of a well-specified prior concentrates at the true index value, i.e.,

$$\pi^{\infty}_{\mu_0} = \delta_{\theta^*} \tag{44}$$

or, equivalently,

$$Support\left(\pi^{\infty}_{\mu_0}\right) = \{\theta^*\}. \tag{45}$$

A seminal result by Berk (1966) generalizes Doob's consistency theorem to the case of not necessarily well-specified priors. To state Berk's Theorem, we need the following definition due to Kullback and Leibler (1951).

Definition 2. The *Kullback-Leibler (KL) divergence* of φ' from φ (also called the *relative entropy of φ with respect to φ'*) is defined as

$$D_{KL}(\varphi||\varphi') = \int_{Support(\varphi)} \frac{d\varphi}{dm}\left[\ln \frac{d\varphi/dm}{d\varphi'/dm}\right] dm \tag{46}$$

$$= E_\varphi\left[\ln \frac{d\varphi}{d\varphi'}\right]. \tag{47}$$

Note that the KL-divergence (46) can take on values in $[0, \infty)$ such that $D_{KL}(\varphi||\varphi') = 0$ if, and only if, $\varphi = \varphi'$. The KL-divergence of φ' from φ gives us thus some notion about "how far φ' is away from φ" without being a fully fledged metric.[12]

[10]Whenever we henceforth speak of *emerging* posteriors or *emerging* priors, the qualification "with probability one" is implicitly included.

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

[12]The KL-divergence is asymmetric and does not satisfy the triangle inequality.

Theorem 0. (Berk 1966). The emerging posterior $\pi_{\mu_0}^{\infty}$ of any prior μ_0 concentrates at the non-empty subset $\Theta_{\mu_0}^{*} \subseteq Support\,(\mu_0)$ consisting of the KL-divergence minimizers φ_θ from the true measure φ_{θ^*}. That is,

$$Support\left(\pi_{\mu_0}^{\infty}\right) \subseteq \Theta_{\mu_0}^{*} \tag{48}$$

such that

$$\Theta_{\mu_0}^{*} = \arg \min_{\theta \in Support(\mu_0)} D_{KL}(\varphi_{\theta^*} || \varphi_\theta). \tag{49}$$

According to Berk's Theorem, Bayesian learning results in an emerging posterior that attaches probability one to measures in the prior's support that minimize the KL-divergence from the true measure. Since the KL-divergence of some φ_θ from the true measure, i.e., $D_{KL}(\varphi_{\theta^*}||\varphi_\theta)$, will play a prominent role in the formal results of this paper (cf. Theorem 2), let us recall the original motivation of Kullback and Leibler (1951) for their divergence concept in terms of "information-based hypothesis discrimination". Let H_θ, $\theta \in \Theta$, be the hypothesis that observation x was drawn in accordance with the probability measure φ_θ so that H_{θ^*} stands for the actually true hypothesis. These authors then define

$$\ln \frac{d\varphi_{\theta^*}(x)}{d\varphi_\theta(x)} \tag{50}$$

"as *the information in x for discrimination between H_{θ^*} and H_θ*" (p. 80, with adjusted notation). Note that (50) is mathematically equivalent[13] to

$$\ln \frac{\pi_{\mu_0}(\theta^* \mid x)}{\pi_{\mu_0}(\theta \mid x)} - \ln \frac{\mu_0(\theta^*)}{\mu_0(\theta)} \tag{51}$$

where $\mu_0(\theta)$ is the decision maker's prior and $\pi_{\mu_0}(\theta \mid x)$ his posterior belief that hypothesis H_θ is the true one. From the difference (51) it is straightforward to see that (50) measures (in logs) how much the information in form of the observed data point x pushes the decision maker towards the true hypothesis (measured by his posterior beliefs) in comparison to from where he started (measured by his prior beliefs). The definition of the KL-divergence $D_{KL}(\varphi_{\theta^*}||\varphi_\theta)$ is then just the "mean information for discrimination between H_{θ^*} and H_θ per observation from φ_{θ^*}" (p. 80, with adjusted notation).

[13]By Bayes' rule, we have that

$$\frac{d\varphi_{\theta^*}(x)}{d\varphi_\theta(x)} = \frac{\pi_{\mu_0}(\theta^* \mid x)/\mu_0(\theta^*)}{\pi_{\mu_0}(\theta \mid x)/\mu_0(\theta)}.$$

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,

$$\{\theta^*\} = \arg \min_{\theta \in Support(\mu_0)} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) \tag{52}$$

so that (48) becomes (45).

4 Bayesian learning with multiple priors

4.1 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 a unique prior μ_0, we now consider a non-empty set of priors $\mathcal{M}_0 \subseteq \triangle^n$. Suppose, for the moment, that the decision maker forms posteriors from all his priors in $\mathcal{M}_0$. Then he will end up with the following set of emerging posteriors after observing an unlimited amount of statistical information

$$\Pi^\infty = \bigcup_{\mu_0 \in \mathcal{M}_0} \left\{ \pi^\infty_{\mu_0} \right\}. \tag{53}$$

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 with probability one at the true measure φ_{θ^*}, i.e.,

$$\Pi^\infty = \{\delta_{\theta^*}\}. \tag{54}$$

For the more general case of not-necessarily well-specified priors, we obtain, by Berk's Theorem 0, that every $\pi^\infty_{\mu_0}$ in (53) has support on some non-empty subset of KL-divergence minimizers (49). In particular, if $\delta_\theta \in \mathcal{M}_0$, then also $\delta_\theta \in \Pi^\infty$. As a consequence, the set of emerging posteriors (53) might become quite implausible if not all priors are well-specified (also recall our argument from Section 2.4).

Example 3. Revisit the "Coin tossing" Example 2. Suppose that the parameter space is given as

$$\Theta = \{0.01, 0.99\} \tag{55}$$

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

$$\Pi^\infty = \{\delta_{0.01}, \delta_{0.99}\} \tag{56}$$

because all mixed priors as well as $\delta_{0.99}$ converge to $\delta_{0.99}$ whereas $\delta_{0.01}$ trivially 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 emerging (=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.,

$$X_1, ..., X_t \longmapsto \mathcal{M}_{0,R}^t,\ t = 1, 2, ..., \tag{57}$$

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

$$\mathcal{M}_{0,R}^t \subseteq \mathcal{M}_0. \tag{58}$$

Note that, by (58), previously rejected priors might reappear later on if they are supported by new data. To illustrate the concept of a prior selection rule, consider the following two (extreme) examples of perceivable rules.

Example 4. 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.,

$$\mathcal{M}_{0,ML}^t = \arg\max_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} \prod_{i=1}^t \frac{d\varphi_\theta}{dm} \cdot \mu_0(\theta). \tag{59}$$

□

Example 5. 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.,

$$\mathcal{M}_{0,MLL}^{t}=\arg\max_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta}\ln\prod_{i=1}^{t}\frac{d\varphi_\theta}{dm}\cdot\mu_0(\theta)\,. \tag{60}$$

□

If $\mathcal{M}_0$ only contains Dirac (i.e., 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 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.

4.2 Emerging priors defined as cluster points

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 priors that almost surely 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 (a.s. P_{θ^*}) *cluster* or the (a.s. P_{θ^*}) *limit points* of sequences $\left\{\mathcal{M}_{0,R}^{t}\right\}_{t\in\mathbb{N}}$ as the priors in $\mathcal{M}_0$ that emerge under the prior selection rule R.

Denote by $\overline{\lim}\mathcal{M}_{0,R}^{t}$ the set that contains all *cluster points* in $\triangle^n$ of the 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 $V\cap\mathcal{M}_{0,R}^{t}\neq\emptyset$.[14] Conversely,

[14]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).

denote by $\underline{\lim}\mathcal{M}_{0,R}^{t}$ 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 $t\geq T$, $V\cap\mathcal{M}_{0,R}^{t}\neq\emptyset$.[15] Whereas every limit point is a cluster point the converse is not true, implying

$$\underline{\lim}\mathcal{M}_{0,R}^{t}\subseteq\overline{\lim}\mathcal{M}_{0,R}^{t}. \tag{61}$$

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

Example 6. 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 $\mathcal{M}_{0}=\{\mu_{0}',\mu_{0}''\}$ with

$$\mu_{0}'=\delta_{\theta_{1}},\ \mu_{0}''=\delta_{\theta_{2}}. \tag{62}$$

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[16] as prior selection rule. Observe that

$$\sum_{\theta\in\Theta}\prod_{i=1}^{t}\frac{d\varphi_{\theta}}{dm}\cdot\mu_{0}'(\theta) \geq \sum_{\theta\in\Theta}\prod_{i=1}^{t}\frac{d\varphi_{\theta}}{dm}\cdot\mu_{0}''(\theta)\Leftrightarrow \tag{63}$$

$$\prod_{k=1}^{t}\frac{d\phi_{\theta_{1}}}{dm}(X_{k}) \geq \prod_{k=1}^{t}\frac{d\phi_{\theta_{2}}}{dm}(X_{k})\Leftrightarrow \tag{64}$$

$$\ln\frac{\prod_{k=1}^{t}\frac{d\phi_{\theta_{1}}}{dm}(X_{k})}{\prod_{k=1}^{t}\frac{d\phi_{\theta_{2}}}{dm}(X_{k})} \geq 0\Leftrightarrow \tag{65}$$

$$\ln\frac{\exp\left[-\frac{1}{2}\sum_{k=1}^{t}(X_{k}-\theta_{1})^{2}\right]}{\exp\left[-\frac{1}{2}\sum_{k=1}^{t}(X_{k}-\theta_{2})^{2}\right]} \geq 0\Leftrightarrow \tag{66}$$

$$\ln\exp\left[(\theta_{1}-\theta_{2})\sum_{k=1}^{t}X_{k}-\frac{t}{2}\left(\theta_{1}^{2}-\theta_{2}^{2}\right)\right] \geq 0\Leftrightarrow \tag{67}$$

$$-2\sum_{k=1}^{t}X_{k} \geq 0 \tag{68}$$

[15] 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).

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

implying that

$$\mathcal{M}_{0,ML}^{t}=\begin{cases}\{\mu_0'\} & \text{if } \sum_{k=1}^{t} X_k<0\\ \{\mu_0',\mu_0''\} & \text{if } \sum_{k=1}^{t} X_k=0\\ \{\mu_0''\} & \text{if } \sum_{k=1}^{t} X_k>0\end{cases} \tag{69}$$

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 (a.s. P_{θ^*})

$$\underline{\lim}\mathcal{M}_{0,ML}^{t}=\emptyset. \tag{70}$$

On the other hand, we (a.s. P_{θ^*}) obtain the non-empty set of cluster points

$$\overline{\lim}\mathcal{M}_{0,ML}^{t}=\{\mu_0',\mu_0''\}\ . \tag{71}$$

□

To take $\{\mu_0',\mu_0''\}$ rather than the empty set as the set of emerging 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-emerging priors.

Definition 3. We define set of *R-emerging priors*, denoted $\mathcal{M}_{0,R}^{\infty}$, as the set of (a.s. P_{θ^*}) cluster points of $\left\{\mathcal{M}_{0,R}^{t}\right\}_{t\in\mathbb{N}}$ in $\mathcal{M}_0$, i.e.,

$$\mathcal{M}_{0,R}^{\infty}\equiv\mathcal{M}_0\cap\overline{\lim}\mathcal{M}_{0,R}^{t}\text{ a.s. }P_{\theta^*}. \tag{72}$$

In words: The set of R-emerging priors consists of all priors in $\mathcal{M}_0$ that will almost surely recur– either as elements or arbitrarily close to elements–in infinitely many sets of the sequence $\left\{\mathcal{M}_{0,R}^{t}\right\}_{t\in\mathbb{N}}$.

4.3 Emerging posteriors and vanishing ambiguity

Fix some prior selection rule R. We stipulate that all emerging posteriors must have been formed from R-emerging priors only.

Definition 4. The set of *emerging posteriors* under the prior selection rule R, denoted Π_R^∞, is defined as

$$\Pi_R^\infty \equiv \bigcup_{\mu_0 \in \mathcal{M}_{0,R}^\infty} \left\{ \pi_{\mu_0}^\infty \right\}. \tag{73}$$

Instead of the above definition, one might alternatively define the set of emerging posteriors as the set of cluster points of the sequence $\{\Pi_R^t\}_{t\in\mathbb{N}}$ that emerge whereby

$$\Pi_R^t \equiv \bigcup_{\mu_0 \in \mathcal{M}_{0,R}^t} \left\{ \pi_{\mu_0} (\cdot \mid X_1, ..., X_t) \right\}. \tag{74}$$

Jan Werner had asked us whether both definitions are equivalent. The following proposition (proved in the Appendix) shows that this equivalence holds for closed sets of priors.

Proposition 2. Let $\mathcal{M}_{0,R}^\infty$ be non-empty. If the set of priors $\mathcal{M}_0$ is closed, then the set of emerging posteriors Π_R^∞ and the set of (a.s. P_{θ^*}) cluster points of $\{\Pi_R^t\}_{t\in\mathbb{N}}$ coincide, i.e.,

$$\bigcup_{\mu_0 \in \mathcal{M}_{0,R}^\infty} \left\{ \pi_{\mu_0}^\infty \right\} = \overline{\lim} \bigcup_{\mu_0 \in \mathcal{M}_{0,R}^t} \left\{ \pi_{\mu_0}^t \right\} \text{ a.s. } P_{\theta^*}. \tag{75}$$

To ensure the equivalence (75), we will henceforth restrict attention to closed sets of priors only.

Definition 5. We say that *ambiguity vanishes* (with probability one) if, and only if, Π_R^∞ is a singleton, i.e.,

$$\Pi_R^\infty = \left\{ \pi_{\mu_0}^\infty \right\}. \tag{76}$$

Suppose, for example, that every prior in $\mathcal{M}_{0,R}^\infty$ has a unique KL-divergence minimizer in its support. Then (73) becomes, by Berk's Theorem, the following collection of Dirac measures

$$\Pi_R^\infty = \bigcup_{\mu_0 \in \mathcal{M}_{0,R}^\infty} \left\{ \delta_{\hat{\theta}} \mid \left\{ \hat{\theta} \right\} = \arg \min_{\theta \in Support(\mu_0)} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) \right\}. \tag{77}$$

In that case, vanishing ambiguity means

$$\Pi_R^\infty = \{\delta_{\hat{\theta}}\} \tag{78}$$

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.

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

$$\mathcal{M}_{0,\alpha}^{t} = \left\{ \mu_0' \in \mathcal{M}_0 \mid \sum_{\theta \in \Theta} \prod_{i=1}^{t} \frac{d\varphi_\theta}{dm} \cdot \mu_0'(\theta) \geq \alpha \cdot \max_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} \prod_{i=1}^{t} \frac{d\varphi_\theta}{dm} \cdot \mu_0(\theta) \right\}. \quad (79)$$

Under some regularity assumptions[17], ES-2007 derive their Claim 3 (p. 1301) according to which all α-emerging 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 7. Revisit the coin tossing situation of Example 2 and suppose that the parameter space is given as

$$\Theta = \{0.49, 0.5, 0.99\} \quad (80)$$

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

[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).

$\mathcal{M}_0 = \{\mu_0', \mu_0''\}$ such that

$$\mu_0' = \varepsilon \cdot \delta_{0.49} + (1-\varepsilon) \cdot \delta_{0.99}, \tag{81}$$

$$\mu_0'' = \delta_{0.5}, \tag{82}$$

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 emerging 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, ES-2007's Theorem 1 (p. 1288) for our one-urn environment.

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

$$\Pi_\alpha^\infty = \{\delta_{\theta^*}\}. \tag{83}$$

5 New results

5.1 The γ-maximum expected loglikelihood rule

Neither does the ES-2007 α-maximum expected likelihood rule allow for non-vanishing STP violations nor does the 'convergence to the truth' result (83) depend on the parameter value α. To tick all the boxes of our wish list for a prior selection rule in Section 2.3, we therefore introduce a new prior selection rule as an alternative to the ES-2007 α-maximum expected likelihood rule.

Definition 7. The *γ-maximum expected loglikelihood rule.* Fix some $\gamma \in [1, \infty)$. The set of emerging priors after observing a sample drawn from $X_1, ..., X_t$ is given as

$$\mathcal{M}_{0,\gamma}^t = \left\{ \mu_0' \in \mathcal{M}_0 \mid \sum_{\theta \in \Theta} \ln \prod_{i=1}^t \frac{d\varphi_\theta}{dm} \cdot \mu_0'(\theta) \geq \gamma \cdot \max_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} \ln \prod_{i=1}^t \frac{d\varphi_\theta}{dm} \cdot \mu_0(\theta) \right\} \tag{84}$$

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

$$\max_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} \ln \prod_{i=1}^{t} \frac{d\varphi_\theta}{dm} \cdot \mu_0(\theta) \leq 0, \tag{85}$$

and

$$\mathcal{M}_{0,\gamma}^t = \mathcal{M}_{0,MLL}^t = \arg \max_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} \ln \prod_{i=1}^{t} \frac{d\varphi_\theta}{dm} \cdot \mu_0(\theta), \tag{86}$$

else.

If (85) holds, the decision maker judges, by (84), priors as plausible whose expected loglikelihood is sufficiently close to the expected loglikelihood of the maximizing prior. The exact meaning of "sufficiently close" is driven by the value of γ whereby greater values of γ will result in greater sets of priors $\mathcal{M}_{0,\gamma}^t$ for fixed values of expected loglikelihoods. We are grateful to Daniele Pennesi who suggested to us to call γ a "stubbornness parameter" since it measures a decision maker's reluctance to give up on priors in the light of new statistical evidence.

To see the intuition behind the formal difference between the ES-2007 α-maximal expected likelihood, on the one hand, and our γ-maximal expected loglikelihood rule, on the other hand, consider the analogy to risk-neutral versus strictly risk averse EU maximization with respect to multiple priors. If likelihoods are taken as *prizes/outcomes/currency*, 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, our decision maker who uses loglikelihoods as utils is more cautious (risk-averse) than the ES-2007 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 the–more cautious–γ-maximal expected loglikelihood rule decision maker.

Remark. Inequality (85) 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 (85) is violated to the effect that the decision maker deals with a strictly positive maximal expected loglikelihood, we simply suppose, by Definition 7, that the γ-maximum expected loglikelihood rule reduces to the maximum expected loglikelihood rule (86).

5.2 Emerging posteriors under the γ-maximum expected loglikelihood rule

The following Theorem constitutes our main formal result.

Theorem 2. The set of posteriors that emerge under the γ-maximum expected loglikelihood rule is given as

$$\Pi_{\gamma}^{\infty} = \bigcup_{\mu_0 \in \mathcal{M}_{0,\gamma}^{\infty}} \left\{ \pi_{\mu_0}^{\infty} \right\} \tag{87}$$

such that

$$\begin{aligned} \mathcal{M}_{0,\gamma}^{\infty} &= \left\{ \mu_0' \in \mathcal{M}_0 \mid \sum_{\theta \in \Theta} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) \cdot \mu_0' (\theta) \leq \min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) \cdot \mu_0 (\theta) \right. \\ & \left. + \left(1 - \frac{1}{\gamma} \right) \sum_{\theta \in \Theta} E_{\varphi_{\theta^*}} \left[- \ln \frac{d\varphi_\theta}{dm} \right] \cdot \mu_0' (\theta) \right\} \end{aligned} \tag{88}$$

whenever the *expected cross-entropy*

$$\sum_{\theta \in \Theta} E_{\varphi_{\theta^*}} \left[- \ln \frac{d\varphi_\theta}{dm} \right] \cdot \mu_0' (\theta) \tag{89}$$

is positive[18], or

$$\mathcal{M}_{0,\gamma}^{\infty} = \arg \min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) \cdot \mu_0 (\theta) \tag{90}$$

else.

Recall from Section 3.2 that the KL-divergence $D_{KL}(\varphi_{\theta^*} || \varphi_\theta)$ can be interpreted as a measure for the "mean information for discrimination under φ_{θ^*}" between the hypothesis that θ^* is the true index value versus the hypothesis that θ is the true value. By Theorem 2, the emerging posteriors under the γ-maximum expected loglikelihood rule are formed from priors that give rise to an expected KL-divergence which is sufficiently close to the minimal expected KL-divergence over all priors in $\mathcal{M}_0$.

The γ-maximum expected loglikelihood rule thus selects (with probability one) priors that put large probability mass on indices which are not strongly discriminated against by information in the form of data generated by the true distribution φ_{θ^*}. The possible interpretation of the γ-maximum expected loglikelihood rule in terms of "information-based discrimination of priors" is, in our opinion, a strong argument in favour of the

[18] A negative expected cross-entropy is impossible for the finite but not for the continuous case.

plausibility of this prior selection rule. In what follows, we further explore how the set of emerging posteriors (87) depends on different values of the stubbornness parameter γ.

5.3 The role of the stubbornness parameter γ

The only index that the γ-maximum expected loglikelihood rule does not discriminate against (with probability one) is the true index θ^*. As a consequence, priors that are well-specified will be particularly rewarded by the γ-maximum expected loglikelihood rule.

Suppose, for example, that the Dirac measure which attaches probability one to the true index θ^* is one of the priors, i.e., $\delta_{\theta^*} \in \mathcal{M}_0$. In this case, we have that

$$\{\delta_{\theta^*}\} = \arg\min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) \cdot \mu_0(\theta)$$

whereby the expected KL-divergence evaluated at δ_{θ^*} becomes zero. Consequently, the set of emerging priors under the γ-maximum expected loglikelihood rule (88) becomes

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

whenever $\delta_{\theta^*} \in \mathcal{M}_0$.

Rewrite now the cross-entropy as the sum of the KL-divergence of φ_θ from φ_{θ^*} and the *Shannon entropy* of φ_{θ^*}, i.e.,

$$E_{\varphi_{\theta^*}} \left[-\ln \frac{d\varphi_\theta}{dm} \right] = D_{KL}(\varphi_{\theta^*} || \varphi_\theta) + E_{\varphi_{\theta^*}} \left[-\ln \frac{d\varphi_{\theta^*}}{dm} \right].$$

Since we always have that

$$\arg\min_{\mu_0 \in \mathcal{M}_0} \sum_{\theta \in \Theta} D_{KL}(\varphi_{\theta^*} || \varphi_\theta) \cdot \mu_0(\theta) \geq 0,$$

a sufficient condition for $\mu_0 \in \mathcal{M}_0$ surviving as an emerging prior in $\mathcal{M}_{0,\gamma}^{\infty}$ is that

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

This gives us immediately the following result about the emerging priors under the γ-maximum expected loglikelihood rule in terms of the stubbornness parameter γ.

Proposition 3. Suppose the Shannon entropy of φ_{θ^*} is strictly positive, i.e.,

$$E_{\varphi_{\theta^*}}\left[-\ln\frac{d\varphi_{\theta^*}}{dm}\right] > 0. \tag{91}$$

Then there exists for every $\mu_0 \in \mathcal{M}_0$ some sufficiently large $\hat{\gamma}$ such that $\mu_0 \in \mathcal{M}_{0,\gamma}^{\infty}$ for all $\gamma \geq \hat{\gamma}$.

The following result about the emerging posteriors under the γ-maximum expected loglikelihood rule follows from a combination of Proposition 3 and Berk's Theorem.

Corollary 1. Suppose that (91) holds.

(i) If

$$\left\{\hat{\theta}\right\} = \arg\min_{\theta\in Support\left(\mu_0'\right)} D_{KL}(\varphi_{\theta^*}||\varphi_{\theta}) \tag{92}$$

for some prior $\mu_0' \in \mathcal{M}_0$, then there exists some sufficiently large $\hat{\gamma}$ such that $\delta_{\hat{\theta}} \in \Pi_{\gamma}^{\infty}$ for all $\gamma \geq \hat{\gamma}$.

(ii) In particular, if $\delta_{\hat{\theta}} \in \mathcal{M}_0$ for any $\hat{\theta} \in \Theta$, then there exists some sufficiently large $\hat{\gamma}$ such that $\delta_{\hat{\theta}} \in \Pi_{\gamma}^{\infty}$ for all $\gamma \geq \hat{\gamma}$.

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

Whereas ambiguity vanishes under the ES-2007 α-maximum expected likelihood rule for arbitrary values of α, ambiguity does not vanish under the γ-maximum expected loglikelihood rule whenever the value of γ is sufficiently large. Consider now the opposite case where γ is as small as possible, i.e., $\gamma = 1$. In this case, the γ-maximum expected loglikelihood rule becomes the maximum expected loglikelihood rule of Example 5, i.e., (88) becomes

$$\mathcal{M}_{0,\gamma=1}^{\infty} = \arg\min_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_{\theta})\cdot\mu_0(\theta). \tag{93}$$

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

Corollary 2. Suppose that $\gamma = 1$. If $\mu_0'(\theta^*) > 0$ for some

$$\mu_0' \in \arg\min_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_{\theta})\cdot\mu_0(\theta), \tag{94}$$

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

$$\Pi_{\gamma=1}^{\infty} = \{\delta_{\theta^*}\}. \tag{95}$$

In particular, we have (95) whenever $\delta_{\theta^*} \in \mathcal{M}_0$.

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 8. Revisit the coin tossing Example 8 where

$$\Theta = \{0.49, 0.5, 0.99\} \tag{96}$$

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

$$\mu_0' = \varepsilon \cdot \delta_{0.49} + (1-\varepsilon) \cdot \delta_{0.99}, \tag{97}$$

$$\mu_0'' = \delta_{0.5}. \tag{98}$$

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

$$\sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_\theta) \cdot \mu_0''(\theta) < \sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_\theta) \cdot \mu_0'(\theta) \tag{99}$$

$$\Leftrightarrow$$

$$E_{\varphi_{\theta^*}}[\ln d\varphi_{\theta^*}] - E_{\varphi_{\theta^*}}[\ln d\varphi_{0.5}] < \varepsilon \cdot \left(E_{\varphi_{\theta^*}}[\ln d\varphi_{\theta^*}] - E_{\varphi_{\theta^*}}[\ln d\varphi_{\theta^*}]\right) + (1-\varepsilon) \cdot \left(E_{\varphi_{\theta^*}}[\ln d\varphi_{\theta^*}] - E_{\varphi_{\theta^*}}[\ln d\varphi_{0.99}]\right) \tag{100}$$

$$\Leftrightarrow$$

$$E_{\varphi_{\theta^*}}[\ln d\varphi_{0.5}] > \varepsilon \cdot E_{\varphi_{\theta^*}}[\ln d\varphi_{\theta^*}] + (1-\varepsilon) \cdot E_{\varphi_{\theta^*}}[\ln d\varphi_{0.99}]$$

since

$$E_{\varphi_{\theta^*}}[\ln d\varphi_{\theta^*}] > E_{\varphi_{\theta^*}}[\ln d\varphi_{0.5}] > E_{\varphi_{\theta^*}}[\ln d\varphi_{0.99}]. \tag{101}$$

To be concrete observe that

$$E_{\varphi_{\theta^*}}[\ln d\varphi_{0.5}] = (1-0.49) \cdot \ln(1-0.5) + 0.49 \cdot \ln 0.5 \approx -0.693147$$

$$E_{\varphi_{\theta^*}}[\ln d\varphi_{\theta^*}] = (1-0.49) \cdot \ln(1-0.49) + 0.49 \cdot \ln 0.49 \approx -0.692947$$

$$E_{\varphi_{\theta^*}}[\ln d\varphi_{0.9999}] = (1-0.49) \cdot \ln(1-0.99) + 0.49 \cdot \ln 0.99 \approx -2.35356$$

so that any $\varepsilon > 0$ satisfying

$$\begin{aligned} \varepsilon &< \frac{E_{\varphi_{\theta^*}}\left[\ln d\varphi_{0.5}\right] - E_{\varphi_{\theta^*}}\left[\ln d\varphi_{0.99}\right]}{E_{\varphi_{\theta^*}}\left[\ln d\varphi_{\theta^*}\right] - E_{\varphi_{\theta^*}}\left[\ln d\varphi_{0.99}\right]} \qquad (102) \\ &\Leftrightarrow \\ \varepsilon &< 0.99988 \qquad (103) \end{aligned}$$

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 (i.e., $\gamma = 1$). However, in the above example the decision maker learns the almost true measure $\varphi_{0.5}$ rather than the true measure $\varphi_{0.49}$. The well-specified prior $\varepsilon \cdot \delta_{0.49} + (1-\varepsilon) \cdot \delta_{0.99}$ is here rejected as implausible because the positive weight on the unlikely parameter $\theta = 0.99$ pulls down the expected loglikelihood of this prior.

6 The Ellsberg one-urn experiment revisited

To show that the γ-maximum expected loglikelihood rule indeed ticks all the boxes of the wish list from Section 2.3, we now revisit the Ellsberg one-urn experiment with increasing statistical information.

Assumption 1. We assume that $\theta^* = 30$; that is, we set $\varphi_{\theta^*}(\omega_2) = \frac{1}{3}$ as the true probability that a black ball will be drawn from the urn.

The following result (proved in the Appendix) combines Proposition 1 with Corollary 1.

Proposition 4. Suppose that Assumption 1 holds. If 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.

To focus our analysis, we next impose the following assumption on the *a priori* decision situation in the Ellsberg one-urn experiment.[19]

Assumption 2. We assume that the set of priors includes all Dirac measures on $(\Theta, \mathcal{F})$, i.e., $\delta_\theta \in \mathcal{M}_0$ for all $\theta \in \Theta$.

Because of Assumption 2 we can easily calculate the value of γ that is needed for ensuring that these degenerate priors also emerge as posteriors (cf. Appendix).

Proposition 5. Suppose that Assumptions 1 and 2 hold. For any $\theta \in \Theta$, we have that

$$\delta_\theta \in \Pi_\gamma^\infty \tag{104}$$

if, and only if,

$$\gamma \geq \frac{\frac{1}{3} \cdot \ln \frac{\theta}{90} + \frac{1}{3} \cdot \ln \left(\frac{60-\theta}{90}\right)}{\frac{2}{3} \cdot \ln \frac{1}{3}}. \tag{105}$$

Obviously, the right side of the inequality (105) 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 (105), identical for parameters θ and $\theta' = 60 - \theta$.

To be concrete, Table 1 lists, for all index values in Θ, the (approximate) values of γ such that (105) holds with equality.

[19] Note that this assumption holds under Gilboa and Schmeidler's (1989, p. 142) 'extreme case'.

$\theta;\theta'$	γ	$\theta;\theta'$	γ
$1;59$	2.24014	$16;44$	1.11178
$2;58$	1.93245	$17;43$	1.09466
$3;57$	1.75583	$18;42$	1.07935
$4;56$	1.63296	$19;41$	1.06571
$5;55$	1.5396	$20;40$	1.05361
$6;54$	1.46497	$21;39$	1.04292
$7;53$	1.40332	$22;38$	1.03357
$8;52$	1.35122	$23;37$	1.02548
$9;51$	1.30645	$24;36$	1.01858
$10;50$	1.26751	$25;35$	1.01282
$11;49$	1.23333	$26;34$	1.00816
$12;48$	1.20311	$27;33$	1.00457
$13;47$	1.17627	$28;32$	1.00203
$14;46$	1.15233	$29;31$	1.00051
$15;45$	1.13093	30	1

Table 1: The impact of γ on 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.[20] Iff

$$1 \leq \gamma < 1.00051, \tag{106}$$

then

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

iff

$$1.00051 \leq \gamma < 1.00203, \tag{108}$$

then

$$\Pi_\gamma^\infty \cap \mathcal{D} = \{\delta_{29}, \delta_{\theta^*}, \delta_{31}\}; \tag{109}$$

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

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

so that there might be priors in $\mathcal{M}_0$, e.g.,

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

with two different KL-divergence minimizers in their support. Emerging posteriors formed from such priors are not necessarily Dirac measures.

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

$$\Pi_\gamma^\infty \cap \mathcal{D} = \{\delta_1, \delta_2, ..., \delta_{59}\}. \tag{110}$$

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. 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 violate the STP in the *a posteriori* decision situation if, and only if, $\gamma \geq 1.00051$ because

$$\delta_{29}, \delta_{31} \in \Pi_\gamma^\infty \tag{111}$$

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

Table 1 illustrates the general relationship that $\gamma' \geq \gamma$ implies $\Pi_\gamma^\infty \subseteq \Pi_{\gamma'}^\infty$ so that Π_γ^∞ expresses *less ambiguity than* $\Pi_{\gamma'}^\infty$ in the sense of set-inclusion of emerging posteriors for a fixed set of priors. The degree of ambiguity that remains after Bayesian learning under the γ-maximum expected loglikelihood rule is thus sensitive with respect to the stubbornness parameter γ.

7 Concluding remarks and outlook

Nicholls et al.'s (2015) report an experiment in which the number of STP violations does not decline in the amount of 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" parameter measuring the decision maker's reluctance to dismiss his priors in the light of statistical evidence.

The possibility of non-vanishing ambiguity is a potentially attractive feature for future 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, the sensitivity of non-vanishing ambiguity with respect to the stubbornness parameter γ might be 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. In dynamic 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-sensitivity 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.

Appendix: Formal proofs

Existence of an emerging posterior. To see that the limit (43) exists, rewrite, for any Θ', $\pi_{\mu_0}(\Theta' \mid X_1, ..., X_t)$ 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

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

whereby convergence happens with P probability one.□□

Proof of Proposition 2. For notational convenience, we write

$$\pi^t_{\mu_0} \equiv \pi_{\mu_0}(\cdot \mid X_1, ..., X_t).$$

Step 1. Suppose that there exists some $\pi^\infty_{\mu_0^*} \in \Pi^\infty_R$ but

$$\pi^\infty_{\mu_0^*} \notin \overline{\lim} \bigcup_{\mu_0 \in \mathcal{M}^t_{0,R}} \left\{\pi^t_{\mu_0}\right\}, \text{ a.s. } P_{\theta^*}. \tag{112}$$

Since $\mu_0^* \in \mathcal{M}^\infty_{0,R}$ and $\mathcal{M}_0$ is compact, there must be (a.s. P_{θ^*}) some subsequence $\{\mu_0^{t_k}\}_{k \in \mathbb{N}}$ such that $\mu_0^{t_k} \in \mathcal{M}^{t_k}_{0,R}$ for all $k = 1, 2, ...$ which converges to μ_0^*. But then $\pi^t_{\mu_0^{t_k}} \in \Pi^t_R$ for all $k = 1, 2, ...$, implying

$$\pi^\infty_{\mu_0^*} \in \overline{\lim} \bigcup_{\mu_0 \in \mathcal{M}^t_{0,R}} \left\{\pi^t_{\mu_0}\right\}, \text{ a.s. } P_{\theta^*}, \tag{113}$$

a contradiction to (112).

Step 2. Next suppose that there exists some

$$\pi^\infty \in \overline{\lim} \bigcup_{\mu_0 \in \mathcal{M}^t_{0,R}} \left\{\pi^t_{\mu_0}\right\}, \text{ a.s. } P_{\theta^*} \tag{114}$$

but

$$\pi^{\infty} \notin \Pi_R^{\infty}. \tag{115}$$

By (114), there must be (a.s. P_{θ^*}) some subsequence $\left\{\pi^{t_k}_{\mu_0^{t_k}}\right\}_{k\in\mathbb{N}}$ such that $\pi^{t_k}_{\mu_0^{t_k}} \in \Pi_R^{t_k}$ for all $k = 1, 2, ...$ which converges to π^{∞}. By compactness of $\mathcal{M}_0$, we can extract a subsequence $\left\{\mu_0^{t_{k'}}\right\}_{k'\in\mathbb{N}}$ from $\left\{\mu_0^{t_k}\right\}_{k\in\mathbb{N}}$ which converges to some $\mu_0^* \in \mathcal{M}_0$. As this μ_0^* is a cluster point of $\left\{\mathcal{M}_{0,R}^t\right\}_{t\in\mathbb{N}}$, we have that $\mu_0^* \in \mathcal{M}_{0,R}^{\infty}$ so that $\pi_{\mu_0^*}^{\infty} \in \Pi_R^{\infty}$. Since $\left\{\pi^{t_{k'}}_{\mu_0^{t_{k'}}}\right\}_{k'\in\mathbb{N}}$ converges to $\pi_{\mu_0^*}^{\infty}$ as well as to π^{∞}, we must have that $\pi_{\mu_0^*}^{\infty} = \pi^{\infty}$, a contradiction to (115).

Collecting Steps 1 and 2 proves the observation.□□

Proof of Theorem 2. Step 1. By Definition 7, the expected loglikelihood maximizing prior(s) will be in $\mathcal{M}_{0,\gamma}^t$ for all t, i.e.,

$$\arg\max_{\mu_0\in\mathcal{M}_0} \sum_{\theta\in\Theta} \ln \prod_{i=1}^{t} \frac{d\varphi_\theta}{dm} \cdot \mu_0(\theta) \subseteq \mathcal{M}_{0,\gamma}^t. \tag{116}$$

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

$$\arg\min_{\mu_0\in\mathcal{M}_0} \sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_\theta) \cdot \mu_0(\theta) \subseteq \mathcal{M}_{0,\gamma}^{\infty} \text{ a.s. } P_{\theta^*}, \tag{117}$$

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

Step 2. Observe that any

$$\mu_0' \in \arg\min_{\mu_0\in\mathcal{M}_0} \sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_\theta) \cdot \mu_0(\theta) \tag{118}$$

belongs to (88) if (88) is non-empty.

Step 3. Suppose now that

$$\mu_0' \in \mathcal{M}_{0,\gamma}^{\infty} \tag{119}$$

but

$$\mu_0' \notin \arg\min_{\mu_0\in\mathcal{M}_0} \sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_\theta) \cdot \mu_0(\theta). \tag{120}$$

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

$$\sum_{\theta\in\Theta}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0'(\theta) \geq \gamma\cdot\max_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0(\theta) \Leftrightarrow \quad (121)$$

$$\sum_{\theta\in\Theta}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0'(\theta) \geq \gamma\cdot\max_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0(\theta) \Rightarrow \quad (122)$$

$$\lim_{t_k\to\infty}\sum_{\theta\in\Theta}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0'(\theta) \geq \lim_{t_k\to\infty}\gamma\cdot\max_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0(\theta). \quad (123)$$

Focus on the l.h.s. term of (123). Because Θ is finite, we can switch the sum and the limit to obtain

$$\lim_{t_k\to\infty}\sum_{\theta\in\Theta}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0'(\theta) = \sum_{\theta\in\Theta}\lim_{t_k\to\infty}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0'(\theta). \quad (124)$$

Turn now to the r.h.s. term of (123). 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

$$f(y_{t_k},\mu_0) \equiv \sum_{\theta\in\Theta}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0(\theta) \quad (125)$$

$$= y_{t_k}\cdot\mu_0 \quad (126)$$

where

$$y_{t_k} = \left(\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_{\theta_1}}{dm},...,\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_{\theta_n}}{dm}\right) \quad (127)$$

and

$$\mu_0 = (\mu_0(\theta_1),...,\mu_0(\theta_n)). \quad (128)$$

Next define the value function of (125) as

$$M(y_{t_k}) = \max_{\mu_0\in\mathcal{M}_0} f(y_{t_k},\mu_0). \quad (129)$$

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)[21] maximum theorem that the value function $M(y_{t_k})$ is continuous. Consequently, if $\lim_{t_k\to\infty} y_{t_k}$ exists, then

$$\lim_{t_k\to\infty}\gamma\cdot M(y_{t_k}) = \gamma\cdot M\left(\lim_{t_k\to\infty}y_{t_k}\right). \quad (130)$$

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

In other words, if

$$\lim_{t_k\to\infty}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm} \tag{131}$$

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

$$\lim_{t_k\to\infty}\gamma\cdot\max_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0(\theta) \tag{132}$$

$$= \gamma\cdot\max_{\mu_0\in\mathcal{M}_0}\lim_{t_k\to\infty}\sum_{\theta\in\Theta}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0(\theta) \tag{133}$$

$$= \gamma\cdot\max_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta}\lim_{t_k\to\infty}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}\cdot\mu_0(\theta). \tag{134}$$

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

$$\ln\frac{d\varphi_\theta}{dm}(X_1),...,\ln\frac{d\varphi_\theta}{dm}(X_n) \tag{135}$$

that

$$\lim_{t_k\to\infty}\frac{1}{t_k}\sum_{i=1}^{t_k}\ln\frac{d\varphi_\theta}{dm}=E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_\theta}{dm}\right]\text{ a.s. }P_{\theta^*} \tag{136}$$

for any θ. By (136) and using (124) and (134), we obtain that (123) is (a.s. P_{θ^*}) equivalent to

$$\sum_{\theta\in\Theta}E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_\theta}{dm}\right]\cdot\mu_0'(\theta) \geq \gamma\cdot\max_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta}E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_\theta}{dm}\right]\cdot\mu_0(\theta)\Leftrightarrow \tag{137}$$

$$\sum_{\theta\in\Theta}E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_\theta}{dm}\right]\cdot\mu_0'(\theta) \geq \gamma\cdot\left(-\min_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta}-E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_\theta}{dm}\right]\cdot\mu_0(\theta)\right) \tag{138}$$

$$\Leftrightarrow$$

$$-\sum_{\theta\in\Theta}E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_\theta}{dm}\right]\cdot\mu_0'(\theta)+\gamma\cdot E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_{\theta^*}}{dm}\right] \tag{139}$$

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

$$\Leftrightarrow$$

$$\gamma\cdot\sum_{\theta\in\Theta}D_{KL}(\varphi_{\theta^*}||\varphi_\theta)d\mu_0'(\theta)-(1-\gamma)\sum_{\theta\in\Theta}E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_\theta}{dm}\right]\cdot\mu_0'(\theta) \tag{140}$$

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

$$\Leftrightarrow$$

$$\sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_\theta)\cdot\mu_0'(\theta) \tag{141}$$
$$\leq \min_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_\theta)d\mu_0(\theta) + \frac{(1-\gamma)}{\gamma}\sum_{\theta\in\Theta} E_{\varphi_{\theta^*}}\left[\ln\frac{d\varphi_\theta}{dm}\right]\cdot\mu_0'(\theta).$$

This proves that

$$\mu_0' \notin \arg\min_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_\theta)\cdot\mu_0(\theta) \tag{142}$$

is in $\mathcal{M}_{0,\gamma}^{\infty}$ if, and only if, μ_0' is in (88).

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

$$\arg\min_{\mu_0\in\mathcal{M}_0}\sum_{\theta\in\Theta} D_{KL}(\varphi_{\theta^*}||\varphi_\theta)\cdot\mu_0(\theta) = \mathcal{M}_{0,\gamma}^{\infty} \tag{143}$$

whenever (88) is empty.

Collecting results proves the proposition.□□

Proof of Proposition 4. Step 1. Consider the *a priori* decision situation. Analogous to the argumentation for Proposition 1, we have that

$$MEU(f_E h, \mathcal{M}_0\circ\Phi) = \frac{1}{3} \text{ and } MEU(g_E h', \mathcal{M}_0\circ\Phi) = \frac{2}{3}. \tag{144}$$

Further, note that

$$\begin{aligned}
MEU(g_E h, \mathcal{M}_0\circ\Phi) &\leq MEU\left(g_E h, \sum_{\theta\in\Theta}\varphi_\theta(\omega)\,\mu_0'(\theta)\right) && (145)\\
&\leq EU\left(g_E h, \sum_{\theta\in\Theta}\varphi_\theta(\omega)\,\delta_{29}\right) && (146)\\
&= EU(g_E h, \varphi_{29}) && (147)\\
&= \frac{29}{90} < \frac{1}{3} && (148)
\end{aligned}$$

as well as

$$\begin{aligned}
MEU(f_E h', \mathcal{M}_0\circ\Phi) &\leq MEU\left(f_E h', \sum_{\theta\in\Theta}\varphi_\theta(\omega)\,\mu_0''(\theta)\right) && (149)\\
&\leq EU\left(f_E h', \sum_{\theta\in\Theta}\varphi_\theta(\omega)\,\delta_{31}\right) && (150)\\
&= EU(f_E h', \varphi_{31}) && (151)\\
&= \frac{1}{3} + \frac{29}{90} < \frac{2}{3}. && (152)
\end{aligned}$$

Consequently, the inequalities (22)-(23) hold.

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

$$MEU\left(f_E h, \Pi_\gamma^\infty \circ \Phi\right) = \frac{1}{3} \text{ and } MEU\left(g_E h', \Pi_\gamma^\infty \circ \Phi\right) = \frac{2}{3}. \tag{153}$$

The specifications of μ_0' and μ_0'' imply, by Corollary 1, for some sufficiently large γ the existence of some

$$\delta_{\theta'}, \delta_{\theta''} \in \Pi_\gamma^\infty \tag{154}$$

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

$$\begin{aligned} MEU\left(g_E h, \Pi_Z^\infty \circ \Phi\right) &\leq EU\left(g_E h, \sum_{\theta \in \Theta} \varphi_\theta(\omega)\, \delta_{\theta'}\right) && (155)\\ &\leq EU\left(g_E h, \sum_{\theta \in \Theta} \varphi_\theta(\omega)\, \delta_{29}\right) && (156)\\ &< \frac{1}{3} && (157) \end{aligned}$$

as well as

$$\begin{aligned} MEU\left(f_E h', \Pi_Z^\infty \circ \Phi\right) &\leq EU\left(f_E h', \sum_{\theta \in \Theta} \varphi_\theta(\omega)\, \delta_{\theta''}\right) && (158)\\ &\leq EU\left(f_E h', \sum_{\theta \in \Theta} \varphi_\theta(\omega)\, \delta_{31}\right) && (159)\\ &< \frac{2}{3}, && (160) \end{aligned}$$

which proves the inequalities (24)-(25).□□

Proof of Proposition 5. By the proof of Theorem 2 (cf., inequality (137) as well as Step 2.), $\mu_0' \in \mathcal{M}_{0,\gamma}^\infty$ if, and only if,

$$\sum_{\theta' \in \Theta} E_{\varphi_{\theta^*}}\left[\ln \frac{d\varphi_{\theta'}}{dm}\right] \cdot \mu_0'(\theta') \geq \gamma \cdot \max_{\mu_0 \in \mathcal{M}_0} \sum_{\theta' \in \Theta} E_{\varphi_{\theta^*}}\left[\ln \frac{d\varphi_{\theta'}}{dm}\right] \cdot \mu_0(\theta'). \tag{161}$$

By Assumption 2, we have, for any $\theta \in \Theta$,

$$\delta_\theta \in \Pi_\gamma^\infty \tag{162}$$

if, and only if,

$$\sum_{\theta' \in \Theta} E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_{\theta'}}{dm} \right] \cdot \delta_\theta \geq \gamma \cdot \sum_{\theta' \in \Theta} E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_{\theta'}}{dm} \right] \cdot \delta_{\theta^*} \tag{163}$$

$$\Leftrightarrow$$

$$E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_\theta}{dm} \right] \geq \gamma \cdot E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_{\theta^*}}{dm} \right] \tag{164}$$

$$\Leftrightarrow$$

$$\frac{E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_\theta}{dm} \right]}{E_{\varphi_{\theta^*}} \left[\ln \frac{d\varphi_{\theta^*}}{dm} \right]} \leq \gamma \tag{165}$$

$$\Leftrightarrow$$

$$\frac{d\varphi_{\theta^*}(\omega_2) \cdot \ln d\varphi_\theta(\omega_2) + \left(\frac{2}{3} - d\varphi_{\theta^*}(\omega_2)\right) \cdot \ln\left(\frac{2}{3} - d\varphi_\theta(\omega_2)\right)}{d\varphi_{\theta^*}(\omega_2) \cdot \ln d\varphi_{\theta^*}(\omega_2) + \left(\frac{2}{3} - d\varphi_{\theta^*}(\omega_2)\right) \cdot \ln\left(\frac{2}{3} - d\varphi_{\theta^*}(\omega_2)\right)} \leq \gamma \tag{166}$$

$$\Leftrightarrow$$

$$\frac{\frac{1}{3} \cdot \ln \frac{\theta}{90} + \frac{1}{3} \cdot \ln\left(\frac{60-\theta}{90}\right)}{\frac{2}{3} \cdot \ln \frac{1}{3}} \leq \gamma, \tag{167}$$

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.

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.

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.

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.

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.

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

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* **50**, 97-115.

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.