On attitude polarization under Bayesian learning with non-additive beliefs*

Alexander Zimper† Alexander Ludwig‡

February 5, 2009

Abstract

Ample psychological evidence suggests that people's learning behavior is often prone to a "myside bias" or "irrational belief persistence" in contrast to learning behavior exclusively based on objective data. In the context of Bayesian learning such a bias may result in diverging posterior beliefs and attitude polarization even if agents receive identical information. Such patterns cannot be explained by the standard model of rational Bayesian learning that implies convergent beliefs. As our key contribution, we therefore develop formal models of Bayesian learning with psychological bias as alternatives to rational Bayesian learning. We derive conditions under which beliefs may diverge in the learning process despite the fact that all agents observe the same arbitrarily large sample, which is drawn from an "objective" i.i.d. process. Furthermore, one of our learning scenarios results in attitude polarization even in the case of common priors. Key to our approach is the assumption of ambiguous beliefs that are formalized as non-additive probability measures arising in Choquet expected utility theory. More precisely, we focus on neo-additive capacities (Chateauneuf et. al. 2007) as a flexible and parsimonious parametrization of departures from additive probability measures. As a specific feature of our approach, our models of Bayesian learning with psychological bias reduce to rational Bayesian learning in the absence of ambiguity.

*We thank Elias Khalil, Robert Östling, Kip Viscusi, Peter Wakker and an anonymous referee for helpful comments and suggestions. Financial support from Economic Research South Africa (ERSA), the German National Research Foundation (DFG) through SFB 504, the State of Baden-Württemberg and the German Insurers Association (GDV) is gratefully acknowledged.

†Department of Economics and Econometrics, University of Johannesburg, PO Box 524, Auckland Park, 2006, South Africa. E-mail: azimper@uj.ac.za

‡Mannheim Research Institute for the Economics of Aging (MEA); Universität Mannheim; L13, 17; 68131 Mannheim; Germany; Email: ludwig@mea.uni-mannheim.de.

Keywords: Non-additive Probability Measures, Choquet Expected Utility Theory, Bayesian Learning, Bounded Rationality

JEL Classification Numbers: C79, D83

1 Introduction

Several studies in the psychological literature demonstrate that people's learning behavior is prone to effects such as "myside bias" or "irrational belief persistence" (cf., e.g., Baron 2008, Chapter 9). For instance, in a famous experiment by Lord, Ross, and Lepper (1979), subjects supporting and opposing capital punishment were exposed to two purported studies, one confirming and one disconfirming their existing beliefs about the deterrent efficacy of the death penalty. Despite the fact that both groups received the same information, their learning behavior resulted in an increased "attitude polarization" in the sense that their respective posterior beliefs, either in favor or against the deterrent efficacy of death penalty, further diverged. Analogous results on diverging posterior beliefs in the face of identical information have earlier been reported by Pitz, Downing, and Reinhold (1967), Pitz (1969) and Chapman (1973) in the context of Bayesian updating of subjective probabilities. In violation of Bayes' update rule the subjects in these experiments formed biased posteriors that supported their original opinions rather than taking into account the evidence. The learning behavior elicited in these experiments cannot be explained by the standard model of rational Bayesian learning according to which differences in agents' prior beliefs must decrease rather than increase whenever the agents receive identical information. In the economics literature, these phenomena have been referred to as a "confirmatory bias" by Rabin and Schrag (1999). That differential interpretation of identical information is of relevance for economic decisions is, e.g., documented in Kandel and Pearson (1995) who provide empirical evidence that news from public announcements are interpreted differently by traders in stock markets. Models of rational Bayesian learning thus apparently ignore relevant aspects of real-life people's learning behavior.

In this paper we present closed-form models of Bayesian learning that allow for the possibility of a "myside bias" as a generalization of a standard rational Bayesian learning model that was introduced to the economics literature by Tonks (1983), Viscusi and O'Connor (1984) and Viscusi (1985). As our point of departure we assume that the paradigm of rational Bayesian learning may only be violated by agents who have ambiguous beliefs. That is, the beliefs of these agents cannot be described by additive probability measures alone but additionally reflect the agent's ambiguity attitudes. The impact of new information on an agent's beliefs is then two-fold. On the one hand, we take into account rational updating based on objective empirical evidence in accordance with our standard model of rational Bayesian learning. On the other hand, however, we also assume existence of a "myside bias" which results in an "irrational" enforcement of the agents' personal attitudes.

In our formal model a decision maker resolves his uncertainty about the "true" parameter value of a Bernoulli trial, e.g., the probability that a given coin turns up *heads*, by some prior belief. In contrast to standard models of Bayesian learning, however, we consider a decision maker who is ambiguous whereby we formally describe ambiguity by non-additive probability measures, i.e., *capacities*, that arise in Choquet Expected Utility (CEU) theory (Schmeidler 1986,1989; Gilboa 1987). CEU theory was originally developed to express ambiguity attitudes of real-life decision makers' behavior who commit paradoxes of the Ellsberg type (Ellsberg 1961) by violating Savage's (1954) sure-thing principle.

In our model, a decision maker's prior estimate of the parameter is then given as the Choquet expected value of possible parameter-values with respect to such ambiguous beliefs. In order to focus our analysis, we further restrict attention to neo-additive capacities in the sense of Chateauneuf, Eichberger and Grant (2007) according to which an agent's non-additive belief about the likelihood of an event is a weighted average of an ambiguous part and an additive part. Neo-additive capacities are particularly attractive because of their parsimonious parametrization of biased beliefs and the direct psychological interpretations of the respective parameters. With respect to the additive part of the neo-additive capacity, we assume that it is described by some distribution of the Beta-distributions family. Under these assumptions, the decision maker's prior belief about the true parameter value is a weighted average of the ambiguous part and the expected value of the Beta-distribution. According to our interpretation, the expected value of this Beta-distribution is the decision maker's best *rational* guess about the "true" value of the parameter. The ambiguous part of his prior belief is relevant whenever the agent lacks absolute confidence in this guess. The associated weighting parameter accordingly reflects the degree of ambiguity (the lack of confidence) expressed by the decision maker with respect to the true parameter. This lack of confidence is resolved in our model by a second parameter that measures the agent's optimistic versus pessimistic personal attitudes with respect to ambiguity.

In a next step we analyze how the decision maker revises his prior belief in light of new information about the outcomes of i.i.d. Bernoulli trials. To this end we consider a decision maker who uses some Bayesian update rule to generate a conditional non-additive probability measure so that his posterior estimate about the parameter is given as the Choquet expected value with respect to this posterior capacity. In the case of non-additive probability measures there exist several perceivable Bayesian update rules expressing different psychological attitudes towards the interpretation of new information (Gilboa and Schmeidler 1993; Sarin and Wakker 1998). As explained below, this "indeterminacy" of update rules is a direct consequence of the violation of the sure thing principle. We then analyze the consequences of three distinct update rules, the so-called *full Bayesian* (Pires 2002; Eichberger, Grant, and Kelsey 2006; Siniscalchi 2006) as well as the *optimistic* and the *pessimistic* update rules (Gilboa and Schmeidler 1993; Sarin and Wakker 1998). An application of these update rules to some prior belief, in which the agent expresses ambiguity, results in a Bayesian learning process that differs from rational Bayesian learning in that convergence to the "true" probabilities of some objective random process will – in general – not emerge. Rather, updating of beliefs reinforces optimistic, respectively pessimistic, attitudes of the agent thereby giving rise to learning behavior with a "myside bias". In this respect, our results resemble those of Rabin and Schrag (1999) who develop a psychologically biased Bayesian updating model that may give rise to a "myside bias"(or "confirmatory bias") because agents are assumed to misread signals conflicting their prior beliefs with some positive probability. In contrast to our approach, their modification of the standard Bayesian learning model by psychological attitudes is, however, rather *ad hoc* and lacks an axiomatic foundation. The update rules that we consider then differ with respect to the strength of reinforcement of such psychological attitudes.

It is important to emphasize that we derive all these results using a closed-form model of Bayesian learning. Key for developing our analytical solutions is the combination of the semi-additive structure of neo-additive capacities with the standard learning model based on the Beta-distribution. In this respect, the parsimonious representation of non-additive beliefs through the neo-additive capacities is not only useful because of the direct psychological interpretations associated with its parametrization but also from an analytical perspective.

Using this Bayesian learning model we then analyze the beliefs of two heterogeneous agents who have some prior beliefs, receive identical information and then update their beliefs according to some Bayesian update rule with psychological bias. Thereby, we differentiate between a *weak* and a *strong* form of myside bias. The weak form of myside bias is characterized by diverging posterior beliefs of the agents under repeated learning with identical information whereby the beliefs may move into the same direction. According to our interpretation the strong form of myside bias is equivalent to attitude polarization in the sense that the posterior estimates of the two agents move into opposite directions under repeated learning with identical information.

To derive our main results we then consider two scenarios: In our first scenario the two agents have different initial beliefs – that is, one agent is more pessimistic than the other – and update their beliefs based on the same information by applying the same update rule. In our second scenario, the two agents receive the same information but apply different update rules. More precisely, we assume that one agent interprets the new information in a rather optimistic way, whereas the other agent interprets the same information in a rather pessimistic way. In both scenarios the resulting posterior beliefs may exhibit the weak as well as the strong form of myside bias. In the first scenario, our results are driven by the initial gap of belief. The intensification of ambiguity attitudes in the course of the learning process than increases this initial gap and thereby reinforces the initial attitudes. The second result is due to the assumption of different learning rules of the agents. While the two agents may have the same prior beliefs, they interpret the same information differentially, which then leads to an increasing gap in beliefs in the course of the learning process.

The remainder of our analysis is structured as follows. In Section 2 we discuss related literature. Section 3 presents our benchmark model of Bayesian learning with non-ambiguous beliefs and Section 4 introduces ambiguous beliefs. Section 5 discusses updating of ambiguous beliefs under the three different update rules considered in this paper. In Section 6 we derive, under the assumption of Bayesian learning, long-run limit estimates that, in general, do not converge to true probabilities. Section 7 then presents our main results on weak and strong myside bias in the form of diverging beliefs and attitude polarization. Finally, Section 8 concludes. All formal proofs are relegated to the Appendix.

2 Related literature

The effects of polarization have received ample attention in the politico-economic literature. A review of this literature is given in Lindqvist and Östling (2008). Examples are Alesina, Baqir and Easterly (1999), Esteban and Ray (2001) and Fernandez and Levy (2008) who study the effects of polarization on government spending, the provision of public goods as well as redistribution. Concepts for the measurement of polarization, which are relevant for much of the empirical work, are developed in Esteban and Ray (1994) and Duclos, Esteban and Ray (2004).

Our work, however, relates to the literature on the determinants of polarization and we focus our following review of the literature accordingly. More precisely, we review learning models that give rise to some form of polarization or myside biases and thereby distinguish between learning models with additive beliefs and models with biased beliefs that are expressed by ambiguity attitudes.

2.1 Learning with additive beliefs

In our learning model agents revise their probability assessments about the parameters of some stochastic process, e.g., about the probability that a given coin turns up heads or tails, by Bayesian updating. Accordingly, agents have some prior beliefs and form posterior beliefs given the relative frequencies observed in the data. In contrast, according to the frequentist approach, agents learn probabilities by simply adopting relative frequencies observed in a given data sample. Within the frequentist approach, divergence of probability assessments of agents cannot occur if the data are drawn from a stationary stochastic process. Against this background, Kurz (1994a,b; 1996) assumes a non-stationary stochastic process and thereby establishes conditions under which agents may not agree about fundamentals in the long run even if they observe the same data sample. However, the application of a frequentist learning rule in a non-stationary environment is not fully consistent because the rationale for agents to apply a frequentist rule for inferring probabilities when the "underlying" probabilities cannot be learnt by this rule is not clear.[1]

While divergence of beliefs can thus not occur within the frequentist framework in a stationary environment, a similar observation holds true within the Bayesian framework when restricted to additive beliefs. Part of our analysis below is based on a specific model of Bayesian learning with additive beliefs according to which the agents' uncertainty with respect to the parameter of a Binomial distribution is described by a Beta distribution. The fact that additive posteriors converge to the same limit belief in this model, however, can be regarded as a special case of more general results on the *consistency* of (additive) Bayesian estimates, in particular Doob's consistency theorem (Doob 1949; for extensions see Breiman, LeCam, and Schwartz 1964; Lijoi, Pruenster, and Walker 2004). Roughly speaking, Doob's consistency theorem states that, for almost all true parameter values, the Bayesian estimate will converge to this value if an agent observes an i.i.d. process.[2] Moreover, for situations in which there are only finitely many possible observations in every period, Freedman (1963) establishes that the Bayesian estimate will converge to the true value whenever this value belongs to the support of the agent's prior.

In light of Doob's consistency result and its extensions it is practically impossible to establish (at least for the case of finitely many observations in every period such

[1] While agents in our model also apply learning rules by which they will not learn "underlying" probabilities, we motivate the application of these rules by psychological and decision-theoretic arguments.

[2] Related to Doob's consistency theorem is Blackwell and Dubins' (1962) convergence theorem for different additive probability measures within the frequentist framework. In particular, Diaconis and Freedman (1986, Theorem 3) establish a formal link between Doob's consistency theorem and Blackwell and Dubins' convergence theorem by basically showing that the Bayesian estimate is consistent if and only if any corresponding conditional probability measures merge in the weak topology.

as heads versus tails) attitude polarization, or even non-converging posteriors, within the framework of Bayesian learners with additive beliefs if all agents observe the same sample information drawn from an i.i.d. process. In order to nevertheless account for the empirical phenomenon of non-converging posteriors or/and attitude polarization, several authors have tried to circumvent these convergence results within the framework of Bayesian learning with additive beliefs. One approach is to restrict attention to the possibility of a short-run bias only, thereby deliberately ignoring long-run convergence (e.g., Brav and Heaton 2002; Dixit and Weibull 2007). Another line of research is to look into the possibility of weakening the i.i.d. assumption of the above framework. E.g., Lewellen and Shanken (2002) consider cases in which the mean of an exogenous dividend process may not be constant over time. Consequently, the agent can never fully learn the objective parameters of the underlying distribution because observed frequencies do not admit any conclusions about objective probabilities even in the long run. Along the same line, Weitzman (2007) considers a non-stationary exogenous stochastic process so that there is no "true" parameter that could be learnt by the agents. Furthermore, within the context of attitude polarization, Kandel and Pearson (1995) and, more recently, Acemoglu, Chernozhukov and Yildiz (2007) consider two agents with different prior-distributions about imprecise signals from an i.i.d. process. Since these different priors imply different interpretation of new information, these authors avoid convergence of both agents' posteriors according to Doob's consistency theorem because these posteriors are effectively formed by observing two different stochastic processes.

2.2 Learning under ambiguity

While the above approaches try – in one way or another – to reconcile the possibility of attitude polarization with Bayesian learning under the assumption of additive beliefs, our approach drops the assumption of additive beliefs altogether. As a consequence, Doob's consistency theorem does generally not apply so that agents' non-additive posteriors may diverge in the long-run despite the fact that they observe the same data drawn from an i.i.d. process. Moreover, our approach may even allow for diverging posteriors and attitude polarization in the case that agents start out with identical priors. This is impossible for models of Bayesian learning with additive beliefs because additivity implies a unique Bayesian update rule.

Related to our approach, Marinacci (1999) studies a learning environment with non-additive beliefs whereby he considers a decision maker who observes an experiment where the outcomes in each trial are identically and independently distributed with respect to the decision-maker's non-additive belief.[3] In this setup, Marinacci derives for (basically convex) capacities laws of large numbers as counterparts to the additive case thereby admitting for the possibility that ambiguity does not vanish in the long-run. While Marinacci's approach may thus be regarded as a frequentist approach towards non-additive probabilities, our approach is a Bayesian one according to which an agent has a subjective prior belief over the whole event space while he uses sample information from an objective process in order to update his subjective belief. In contrast to our approach the learning behavior of different agents in Marinacci's model must converge to the same limit if they have identical priors. As a consequence there cannot occur attitude polarization within Marinacci's framework under the assumption of common priors.

Epstein and Schneider (2007) also consider a model of learning under ambiguity which shares with our learning model the feature that ambiguity does not necessarily vanish in the long run. Their learning model is based on the *recursive multiple priors* approach (Epstein and Wang 1994; Epstein and Schneider 2003) that restricts conditional *max min expected utility* (MMEU) preferences of Gilboa and Schmeidler (1989) such that dynamic consistency is satisfied. While MMEU theory is closely related to CEU theory restricted to *convex* capacities (e.g., neo-additive capacities for which the degree of optimism is zero), the similarity between Epstein and Schneider's approach and our learning model ends here. As one main difference, the restriction of Epstein and Schneider's approach to dynamically consistent preferences excludes preferences that violate Savage's sure-thing principle as elicited in Ellsberg paradoxes (Ellsberg 1961).[4] Since our learning model does not exclude dynamically inconsistent decision behavior, it can accommodate a broader notion of ambiguity attitudes than the Epstein-Schneider approach, including ambiguity attitudes that are not compatible with the sure-thing principle. Furthermore, Epstein and Schneider establish long-run ambiguity, i.e., the existence of multiple posteriors, under the assumption that the decision-maker permanently receives ambiguous signals, which they formalize via a multitude of different likelihood functions at each information stage in addition to the existence of multiple priors.[5] This introduction of multiple likelihoods is rather ad hoc and it would be interesting to see an axiomatic

[3]Notice that there are several perceivable definitions of independence for capacities. Very loosely speaking, in the context of conditional capacities Marinacci's notion of independence corresponds to the optimistic update rule, ensuring that $\nu(A \mid B) = \nu(A)$ if A and B are independent with respect to the capacity ν.

[4]More precisely, dynamically consistent preferences are at odds with Ellsberg-like behavior whenever the principles of *consequentialism* and *reduction* are satisfied, which is the case in Epstein and Schneider's as well as in our approach (cf. Siniscalchi 2006 and Observation 2 in Zimper 2008).

[5]In the case of learning from ambiguous urns without multiple likelihoods, ambiguity obviously vanishes in the learning process, see Marinacci (2002) for a formal result.

or/and psychological foundation of this approach which goes beyond the mere technical property that multiple likelihoods can sustain long-run ambiguity in the recursive multiple priors framework. On the contrary, our – comparably simple – axiomatically founded model of Bayesian learning with psychological attitudes offers a rather straightforward rationale for why long-run beliefs may be biased even in the case that the decision-maker receives signals that are not ambiguous.

3 The benchmark case: Rational Bayesian learning

In this section we describe in detail a closed-form learning model with additive beliefs as introduced to the economics literature by Viscusi and O'Connor (1984) and Viscusi (1985). Consider the situation of an agent who is uncertain about the probability of an outcome, H, but can observe a statistical experiment with n independent trials where H, resp. T, is a possible outcome that occurs with identical probability.

Formally, we consider a probability space $(\mu, \Omega, \mathcal{F})$ where μ denotes a subjective additive probability measure defined on the events in the event space $\mathcal{F}$. As state space we assume

$$\Omega = \Pi \times S^{\infty},$$

with generic element $\omega = (\pi, s^{\infty})$, where the parameter-space $\Pi = [0, 1]$ collects all possible values of the "true" probability of H in any given trial and the sample-space $S^{\infty} = \times_{i=1}^{\infty} \{H, T\}$ collects all possible (infinite) sequences of outcomes.

We now construct an event space that is rich enough to express our assumptions that after any given number n of trials (i) the agent knows the outcome of each of these trials whereas (ii) he does not make any observations about the true parameter value. These assumptions are central to Bayesian learning. The event space $\mathcal{F}$ is accordingly defined as follows. Endow Π with the Euclidean metric and denote by $\mathcal{B}$ the Borel σ-algebra in Π, i.e., the smallest σ-algebra containing all open subsets of the Euclidean unit interval. Similarly, endow S^{∞} with the discrete metric and denote by $\mathcal{S}^{\infty}$ the Borel σ-algebra in S^{∞}; that is, $\mathcal{S}^{\infty}$ collects all subsets of S^{∞}. Our event space $\mathcal{F}$ is then defined as the standard product algebra of $\mathcal{B}$ and $\mathcal{S}^{\infty}$.

Define $\tilde{\pi} : \Omega \to [0, 1]$ such that $\tilde{\pi}(\pi, s^{\infty}) = \pi$ is the random variable that assigns to every state of the world the true probability of outcome H. For our purpose it is convenient to denote by $\boldsymbol{\pi}$ the event in $\mathcal{F}$ such that $\pi \in \Pi$ is the true probability, i.e.,

$$\boldsymbol{\pi} = \{\omega \in \Omega \mid \tilde{\pi}(\omega) = \pi\}.$$

We assume that the agent's prior over $\tilde{\pi}$ is given as a Beta distribution with parameters $\alpha, \beta > 0$ so that

$$\mu(\boldsymbol{\pi}) = K_{\alpha,\beta}\pi^{\alpha-1}(1-\pi)^{\beta-1}$$

where $K_{\alpha,\beta} = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}$ is a normalizing constant whereby the gamma function is defined as $\Gamma(y) = \int_0^\infty x^{y-1}e^{-x}dx$ for $y > 0$. Let $X_n : \Omega \to \{0,1\}$, with $n = 1, 2, ...$, denote the random variable that takes on value one if a head occurs in the n-th trial and zero otherwise. We assume that, conditional on the parameter-value π, each X_n is Bernoulli distributed with probabilities

$$\mu(\{\omega \in \Omega \mid X_n(\omega) = x\} \mid \boldsymbol{\pi}) = \pi^x (1-\pi)^{1-x} \text{ for } x \in \{0,1\}.$$

Furthermore, denote by $I_n : \Omega \to \{0, ..., n\}$ the random variable counting the number of heads occurring in n trials, i.e., $I_n = \sum_{k=1}^n X_k$. For convenience, we will denote by $\mathbf{I}_n^k$ the event in $\mathcal{F}$ such that outcome H has occurred k-times in the n first trials, i.e.,

$$\mathbf{I}_n^k = \{\omega \in \Omega \mid I_n(\omega) = k\}.$$

Since the X_n are i.i.d. Bernoulli distributed, each I_n is, conditional on the parameter-value π, binomially distributed with probabilities

$$\mu\left(\mathbf{I}_n^k \mid \boldsymbol{\pi}\right) = \binom{n}{k}\pi^k (1-\pi)^{n-k} \text{ for } k \in \{0, ..., n\}.$$

By Bayes' rule we obtain the following posterior probability that π is the true value conditional on information $\mathbf{I}_n^k$

$$\begin{aligned}
\mu\left(\boldsymbol{\pi} \mid \mathbf{I}_n^k\right) &= \frac{\mu\left(\boldsymbol{\pi} \cap \mathbf{I}_n^k\right)}{\mu\left(\mathbf{I}_n^k\right)} \\
&= \frac{\mu\left(\mathbf{I}_n^k \mid \boldsymbol{\pi}\right)\mu(\boldsymbol{\pi})}{\int_{[0,1]} \mu\left(\mathbf{I}_n^k \mid \boldsymbol{\pi}\right)\mu(\boldsymbol{\pi})\, d\pi} \\
&= K_{\alpha+k,\beta+n-k}\pi^{\alpha+k-1}(1-\pi)^{\beta+n-k-1}.
\end{aligned}$$

Example. Suppose that the agent's prior over $\tilde{\pi}$ is given as the *uniform distribution*, which is a Beta distribution with parameters $\alpha = \beta = 1$. Roughly speaking, such an agent would perceive every probability in Π as equally likely. Conditional on information $\mathbf{I}_n^k$, the agent's posterior over $\tilde{\pi}$ is then given by

$$\begin{aligned}
\mu\left(\boldsymbol{\pi} \mid \mathbf{I}_n^k\right) &= \frac{\Gamma(2+n)}{\Gamma(1+k)\Gamma(1+n-k)}\pi^k (1-\pi)^{n-k} \\
&= \frac{(n+1)!}{k!\,(n-k)!}\pi^k (1-\pi)^{n-k},
\end{aligned}$$

whereby the last equality follows from the fact that $\Gamma(y) = (y-1)!$ for $y = 1, 2, \ldots$, (cf. Theorem 8.18 in Rudin 1976).□

The agent's prior estimate of the true probability of H is defined as the expected value of $\tilde{\pi}$ with respect to the prior distribution, i.e., $E[\tilde{\pi}, \mu]$. Accordingly, the agent's posterior estimate of π conditional on information $\mathbf{I}_n^k$ is defined as the expected value of $\tilde{\pi}$ with respect to the resulting posterior distribution, i.e., $E\left[\tilde{\pi}, \mu\left(\cdot \mid \mathbf{I}_n^k\right)\right]$. In the case of a Beta prior we therefore obtain as prior estimate the expected value of the Beta distribution, implying $E[\tilde{\pi}, \mu] = \frac{\alpha}{\alpha+\beta}$. Furthermore, because a prior of the Beta distribution family is a conjugate prior, the agent's posterior $\mu\left(\cdot \mid \mathbf{I}_n^k\right)$ is itself a Beta distribution with parameters $\alpha+k, \beta+n-k$. We therefore have $E\left[\tilde{\pi}, \mu\left(\cdot \mid \mathbf{I}_n^k\right)\right] = \frac{\alpha+k}{\alpha+\beta+n}$ as the agent's posterior estimate of π conditional on information $\mathbf{I}_n^k$. Or, equivalently,

$$E\left[\tilde{\pi}, \mu\left(\cdot \mid \mathbf{I}_n^k\right)\right] = \left(\frac{\alpha+\beta}{\alpha+\beta+n}\right) E[\tilde{\pi}, \mu] + \left(\frac{n}{\alpha+\beta+n}\right) \frac{k}{n} \tag{1}$$

where $\frac{k}{n}$ is the sample mean. That is, the agent's posterior estimate of the probability of H is a weighted average of his prior estimate and the sample mean whereby the weight attached to the sample mean increases in the number of trials.[6] Let π^* denote the "true" probability of outcome H. If the number of trials approaches infinity, i.e., $n \rightarrow \infty$, the sample mean information $\mathbf{I}_n^k$ converges in probability to the sample information $\mathbf{I}^*$ according to which outcome H has occurred with relative frequency π^*. That is, for every $c > 0$, $\lim_{n\rightarrow\infty} prob\left(\left|\mathbf{I}_n^k - \pi^*\right| \leq c\right) = 1$. As a consequence, we obtain the following consistency result for this specific model of Bayesian learning with additive beliefs.

Observation 1: *The posterior estimates $E\left[\tilde{\pi}, \mu\left(\cdot \mid \mathbf{I}_n^k\right)\right]$ of the probability of outcome H converge in probability to the true probability π^* as n approaches infinity.*

Apparently, this closed-form model of rational Bayesian learning cannot account for the learning behavior of agents whose posterior beliefs systematically diverge while they receive the same sample information drawn from an i.i.d. process. By Doob's consistency theorem a similar observation applies to general models of Bayesian learning whenever beliefs are given as additive probability measures.

[6] Tonks (1983) introduces a similar model of rational Bayesian learning in which the agent has a normally distributed prior over the mean of some normal distribution and receives normally distributed information.

4 Ambiguous beliefs

We assume that individuals exhibit *ambiguity attitudes* in the sense of Schmeidler (1989) and who may thus, for example, commit paradoxes of the Ellsberg type (Ellsberg 1961). Following Schmeidler (1989) and Gilboa (1987), we describe such individuals as Choquet Expected Utility (CEU) decision makers, that is, they maximize expected utility with respect to *non-additive beliefs*. Properties of non-additive beliefs are used in the literature for formal definitions of, e.g., ambiguity and uncertainty attitudes (Schmeidler 1989; Epstein 1999; Ghirardato and Marinacci 2002), pessimism and optimism (Eichberger and Kelsey 1999; Wakker 2001; Chateauneuf, Eichberger, and Grant 2007), as well as sensitivity to changes in likelihood (Wakker 2004).[7] The Choquet expected value of a bounded random variable $Y : \Omega \rightarrow \mathbb{R}$ with respect to capacity ν is formally defined as the following Riemann integral extended to domain Ω (Schmeidler 1986):

$$E[Y,\nu] = \int_{-\infty}^{0} \left(\nu\left(\{\omega \in \Omega \mid Y(\omega) \geq z\}\right) - 1\right) dz + \int_{0}^{+\infty} \nu\{\omega \in \Omega \mid Y(\omega) \geq z\}\, dz. \quad (2)$$

Our own approach focuses on non-additive beliefs that are defined as *neo-additive capacities* in the sense of Chateauneuf, Eichberger and Grant (2007).

Definition (Neo-additive Capacities). *For a given measurable space* $(\Omega, \mathcal{F})$ *the neo-additive capacity,* ν*, is defined, for some* $\delta, \lambda \in [0,1]$ *by*

$$\nu(A) = \delta \cdot (\lambda \cdot \omega^{o}(A) + (1-\lambda) \cdot \omega^{p}(A)) + (1-\delta) \cdot \mu(A) \quad (3)$$

for all $A \in \mathcal{F}$ *such that* μ *is some additive probability measure and we have for the non-additive capacities* ω^{o}

$$\begin{aligned} \omega^{o}(A) &= 1 \text{ if } A \neq \emptyset \\ \omega^{o}(A) &= 0 \text{ if } A = \emptyset \end{aligned}$$

and ω^{p} *respectively*

$$\begin{aligned} \omega^{p}(A) &= 0 \text{ if } A \neq \Omega \\ \omega^{p}(A) &= 1 \text{ if } A = \Omega. \end{aligned}$$

The following observation extends a result (Lemma 3.1) of Chateauneuf, Eichberger, and Grant (2007) for finite random variables to the more general case of random variables with a closed and bounded support.

[7] See Wakker (2009) for a textbook treatment.

Observation 2: *Let Y be a closed and bounded random variable. Then the Choquet expected value (2) of Y with respect to a neo-additive capacity (3) is given by*

$$E[Y,\nu] = \delta(\lambda \max Y + (1-\lambda)\min Y) + (1-\delta) E[Y,\mu]. \tag{4}$$

Neo-additive capacities can be interpreted as non-additive beliefs that stand for deviations from additive beliefs such that a parameter δ, the *degree of ambiguity*, measures the lack of confidence the decision maker has in some subjective additive probability distribution μ. Obviously, if there is no ambiguity, i.e., $\delta = 0$, (4) reduces to the standard subjective expected utility representation of Savage (1954). In case there is some ambiguity, however, the second parameter λ measures how much weight the decision maker puts on the best possible outcome of Y when resolving his ambiguity. Conversely, $(1-\lambda)$ is the weight he puts on the worst possible outcome of Y. As a consequence, we interpret λ as the *degree of optimism under ambiguity* whereby $\lambda = 1$, resp. $\lambda = 0$, corresponds to extreme optimism, resp. extreme pessimism, with respect to resolving ambiguity in the decision maker's belief.

Finally, observe that for non-degenerate events, i.e., $A \notin \{\emptyset, \Omega\}$, the neo-additive capacity ν in (3), simplifies to

$$\nu(A) = \delta \cdot \lambda + (1-\delta) \cdot \mu(A). \tag{5}$$

5 Updating ambiguous beliefs

CEU theory has been developed in order to accommodate paradoxes of the Ellsberg type which show that real-life decision-makers violate Savage's *sure-thing principle* according to which preferences over acts shall be unaffected by consequences in states in which the two acts have the same outcome. In this section we demonstrate that abandoning the sure-thing principle bears two important implications for conditional CEU preferences over Savage-acts. First, in contrast to Bayesian updating of additive probability measures, there exist several perceivable Bayesian update rules for non-additive probability measures (cf. Gilboa and Schmeidler 1993, Sarin and Wakker 1998, Pires 2002, Eichberger, Grant and Kelsey 2006, Siniscalchi 2006). Second, any preferences that (strictly) violate the sure-thing principle cannot be updated in a dynamically consistent way. That is, there does not exist any updating rule for capacities such that ex-ante CEU preferences that (strictly) violate the sure-thing principle are updated in a dynamically consistent manner to ex-post CEU preferences.

To see this define the Savage-act $f_B h : \Omega \to X$ such that

$$f_B h(\omega) = \begin{cases} f(\omega) & \text{for } \omega \in B \\ h(\omega) & \text{for } \omega \in \neg B \end{cases}$$

where B is some non-empty event. That is, the act $f_B h$ gives the same consequences as the act f in all states belonging to event B and it gives the same consequences as the act h in all states outside of event B. Recall that Savage's sure-thing principle states that, for all acts f, g, h, h' and all events $B \in \mathcal{F}$,

$$f_B h \succeq g_B h \text{ implies } \; f_B h' \succeq g_B h'. \tag{6}$$

That is, preferences over Savage-acts f and g should be unaffected by any states in which these acts give the same consequences. Let us now interpret event B as new information received by the agent. The sure-thing principle then implies a straightforward way for deriving ex-post preferences $\succeq_B$, conditional on the new information B, from the agent's original preferences $\succeq$ over Savage-acts. Namely, we have

$$f \succeq_B g \text{ if and only if } f_B h \succeq g_B h \text{ for any } h, \tag{7}$$

so that an agent's ex-post preferences over two acts in light of new information B are given as the agent's ex-ante preferences over these acts whenever both acts give arbitrary but identical consequences in states of the world that do not belong to B (i.e., in states that will be declared *impossible* by the new information B). Equation (7) implies for a subjective EU decision-maker

$$f \succeq_B g \Leftrightarrow E[u(f), \mu(\cdot \mid B)] \geq E[u(g), \mu(\cdot \mid B)]$$

where $u : X \to \mathbb{R}$ is a von Neumann-Morgenstern utility function and $\mu(\cdot \mid B)$ is a conditional additive probability measure defined, for all $A, B \in \mathcal{F}$ such that $\mu(B) > 0$, by

$$\mu(A \mid B) = \frac{\mu(A \cap B)}{\mu(B)}.$$

In case the sure-thing principle does not hold, the specification of act h in (7) is no longer arbitrary. For CEU preferences there therefore exist several possibilities of deriving ex post preferences from ex ante preferences. That is, in a CEU framework there exist several perceivable ways of defining a conditional capacity $\nu(\cdot \mid B)$ such that

$$f \succeq_B g \Leftrightarrow E[u(f), \nu(\cdot \mid B)] \geq E[u(g), \nu(\cdot \mid B)].$$

Let us at first consider conditional CEU preferences satisfying, for all acts f, g,

$$f \succeq_B g \text{ if and only if } f_B h \succeq g_B h$$

where h is the so-called conditional certainty equivalent of g, i.e., given information B the agent is indifferent between the act g and the act h that gives in every state of B the same consequence. The corresponding Bayesian update rule for the non-additive beliefs of a CEU decision maker is the so-called full Bayesian update rule which is given by (Eichberger, Grant, and Kelsey 2006)

$$\nu^{FB}(A \mid B) = \frac{\nu(A \cap B)}{\nu(A \cap B) + 1 - \nu(A \cup \neg B)}, \tag{8}$$

where $\nu^{FB}(A \mid B)$ denotes the conditional capacity for event $A \in \mathcal{F}$ given information $B \in \mathcal{F}$.

Observation 3: *An application of the full Bayesian update rule (8) to a prior belief (5) results in the posterior belief=0*

$$\nu^{FB}(A \mid B) = \delta_B^{FB} \cdot \lambda + \left(1 - \delta_B^{FB}\right) \cdot \mu(A \mid B) \tag{9}$$

whereby

$$\delta_B^{FB} = \frac{\delta}{\delta + (1 - \delta) \cdot \mu(B)}. \tag{10}$$

In addition to the full Bayesian update rule we also consider so-called *h-Bayesian update rules* for preferences $\succeq$ over Savage acts as introduced by Gilboa and Schmeidler (1993). That is, we consider some collection of conditional preference orderings, $\left\{\succeq_B^h\right\}$ for all events B, such that for all acts f, g

$$f \succeq_B^h g \text{ if and only if } f_B h \succeq g_B h \tag{11}$$

where

$$h = (x^*, E; x_*, \neg E), \tag{12}$$

with x^* denoting the best and x_* denoting the worst consequence possible and $E \in \mathcal{F}$. For the so-called *optimistic* update rule h is the constant act where $E = \emptyset$. That is, under the optimistic update rule the null-event, $\neg B$, becomes associated with the worst consequence possible. Gilboa and Schmeidler (1993) offer the following psychological motivation for this update rule:

> "[...] when comparing two actions given a certain event B, the decision maker implicitly assumes that had B not occurred, the worst possible outcome [...] would have resulted. In other words, the behavior given B [...] exhibits 'happiness' that B has occurred; the decisions are made as if we are always in 'the best of all possible worlds'."

As corresponding optimistic Bayesian update rule for conditional beliefs of CEU decision makers we obtain

$$\nu^{opt}(A \mid B) = \frac{\nu(A \cap B)}{\nu(B)}. \tag{13}$$

Observation 4: *An application of the optimistic update rule (13) to a prior belief (5) such that*

$$\text{NOT } (\delta = 1 \text{ AND } \lambda = 0) \tag{14}$$

results in the conditional belief

$$\nu^{opt}(A \mid B) = \delta_B^{opt} + \left(1 - \delta_B^{opt}\right) \cdot \mu(A \mid B)$$

with

$$\delta_B^{opt} = \frac{\delta \cdot \lambda}{\delta \cdot \lambda + (1 - \delta) \cdot \mu(B)}.$$

For the *pessimistic* (or Dempster-Shafer) update rule h is the constant act where $E = \Omega$, associating with the null-event,$\neg B$, the best consequence possible. The psychological interpretation for this update rule according to Gilboa and Schmeidler (1993) is as follows:

> "[...] we consider a 'pessimistic' decision maker, whose choices reveal the hidden assumption that all the impossible worlds are the best conceivable ones."

The corresponding pessimistic Bayesian update rule for CEU decision makers is

$$\nu^{pess}(A \mid B) = \frac{\nu(A \cup \neg B) - \nu(\neg B)}{1 - \nu(\neg B)}. \tag{15}$$

Observation 5: *An application of the pessimistic update rule (15) to a prior belief (5) such that*

$$\text{NOT } (\delta = 1 \text{ AND } \lambda = 1) \tag{16}$$

results in the conditional belief

$$\nu^{pess}(A \mid B) = \left(1 - \delta_B^{pess}\right) \cdot \mu(A \mid B)$$

with

$$\delta_B^{pess} = \frac{\delta \cdot (1 - \lambda)}{\delta \cdot (1 - \lambda) + (1 - \delta) \cdot \mu(B)}.$$

Remark. Observe that the conditions (14) and (16) are consistency conditions which ensure that the denominator in the according conditional capacity is not zero so that the conditional capacities are well-defined. In the remainder of the paper we will assume that (14) and (16) hold without explicitly mentioning it. To see the intuition behind these consistency conditions notice that (14), resp. (16), states that extremely pessimistic, resp. optimistic, priors should not be updated by the optimistic, resp. pessimistic, rule.

6 Learning with ambiguous beliefs

In this section we formally link the updating of ambiguous beliefs to Bayesian learning behavior. As a generalization of the Bayesian learning model discussed in Section 3, we consider now a neo-additive prior about the unknown parameter π such that

$$\nu(\boldsymbol{\pi}) = \delta\lambda + (1-\delta)\cdot K_{\alpha,\beta}\pi^{\alpha-1}(1-\pi)^{\beta-1}, \tag{17}$$

i.e., the additive part of this prior is the Beta distribution with parameters α and β. Accordingly, the agent's prior estimate of the true value of π is now given as the Choquet expected value of $\tilde{\pi}$ with respect to his neo-additive prior, i.e.,

$$\begin{aligned} E[\tilde{\pi},\nu] &= \delta(\lambda\max\tilde{\pi} + (1-\lambda)\min\tilde{\pi}) + (1-\delta)E[\tilde{\pi},\mu] \\ &= \delta\lambda + (1-\delta)E[\tilde{\pi},\mu] \end{aligned}$$

by observation 2 and the fact that $\tilde{\pi}$ takes on all values in $[0,1]$.

In the limit of a Bayesian learning process the agent's posterior estimates of π will then converge to $E[\tilde{\pi},\nu(\cdot \mid \mathbf{I}^*)]$, whose value depends on the applied Bayesian update rule. Based on a lemma that we present and prove in the appendix, the following observation characterizes these limit estimates.

Observation 6: *Contingent on the applied update rule – full Bayesian, optimistic, or pessimistic – the agent's estimates about the probability of outcome H converge in probability to the following posterior estimates as n approaches infinity.*

(i) *Full Bayesian learning.*

$$E\left[\tilde{\pi},\nu^{FB}(\cdot \mid \mathbf{I}^*)\right] = \delta^{FB}_{\mathbf{I}^*}\lambda + \left(1-\delta^{FB}_{\mathbf{I}^*}\right)\cdot\pi^*$$

whereby

$$\delta_{\mathbf{I}^*}^{FB} = \frac{\delta}{\delta + (1-\delta) \cdot \mu (\mathbf{I}^*)}.$$

(ii) *Optimistic Bayesian learning.*

$$E \left[\tilde{\pi}, \nu^{opt} (\cdot \mid \mathbf{I}^*)\right] = \delta_{\mathbf{I}^*}^{opt} + \left(1 - \delta_{\mathbf{I}^*}^{opt}\right) \cdot \pi^*$$

whereby

$$\delta_{\mathbf{I}^*}^{opt} = \frac{\delta \cdot \lambda}{\delta \cdot \lambda + (1-\delta) \cdot \mu (\mathbf{I}^*)}.$$

(iii) *Pessimistic Bayesian learning.*

$$E \left[\tilde{\pi}, \nu^{pess} (\cdot \mid \mathbf{I}^*)\right] = (1 - \delta_{\mathbf{I}^*}^{pess}) \cdot \pi^*$$

whereby

$$\delta_{\mathbf{I}^*}^{pess} = \frac{\delta \cdot (1-\lambda)}{\delta \cdot (1-\lambda) + (1-\delta) \cdot \mu (\mathbf{I}^*)}.$$

Observe that the agent's limit estimates are given as a weighted average between the true parameter value π^* and the numbers λ (for full Bayesian learning), 1 (for optimistic Bayesian learning), and 0 (for pessimistic Bayesian learning), respectively. Thus, in contrast to the benchmark case of rational Bayesian learning, the above learning rules will in general not converge to the true parameter value whenever there is some positive initial degree of ambiguity, i.e., $\delta > 0$. More specifically, observe that for all learning rules the limit estimate's weight on the true parameter value decreases in the prior additive probability attached to the limit information so that, in general, the estimate $E [\tilde{\pi}, \nu (\cdot \mid \mathbf{I}^*)]$ is the more biased the smaller is $\mu (\mathbf{I}^*)$. As an intuitive explanation of this formal relationship we could say that the agent's learning behavior puts the less weight on the statistical data the more he is surprised to actually observe this data.

In the light of this interpretation, also notice that the size of $\mu (\mathbf{I}^*)$ depends on the parameters α and β as well as on the accuracy of information received by the agent. For instance, if the agent's prior is given by the uniform distribution, i.e., $\alpha = \beta = 1$, and his information about the sample mean is always *precise* in the sense that $\mathbf{I}_n^k = \frac{k}{n}$ for all n, we encounter the extreme case $\mu (\mathbf{I}^*) = 0$ regardless of the specific value of π^* (cf. Ludwig and Zimper 2008a). The results of this paper, however, are all valid for the general case $0 \leq \mu (\mathbf{I}^*) < 1$, which might arise in the case of imprecise information where $\mathbf{I}_n^k = \left[\frac{k}{n} - \varepsilon, \frac{k}{n} + \varepsilon\right]$ for all n with $\varepsilon > 0$.

Consider now the situation that different learners start out with identical neo-additive priors. The following result formally confirms our intuition that a pessimistic learner will always end up with a smaller posterior estimate than a full Bayesian learner who in turn always ends up with a smaller posterior estimate than an optimistic learner. Furthermore, while an optimistic (pessimistic) learner will always overestimate (underestimate) the true probability of the i.i.d. process, a full Bayesian learner will overestimate (underestimate) this true probability if and only if it is smaller (greater) than his original degree of optimism λ.

Observation 7: *Suppose that $\delta > 0$ and $\lambda \in (0,1)$. Then*

$$E\left[\tilde{\pi}, \nu^{pess}\left(\cdot \mid \mathbf{I}^{*}\right)\right] < E\left[\tilde{\pi}, \nu^{FB}\left(\cdot \mid \mathbf{I}^{*}\right)\right] < E\left[\tilde{\pi}, \nu^{opt}\left(\cdot \mid \mathbf{I}^{*}\right)\right].$$

Moreover, with respect to any "true" probability $\pi^{*} \in (0,1)$ we have for these limit estimates

$$E\left[\tilde{\pi}, \nu^{pess}\left(\cdot \mid \mathbf{I}^{*}\right)\right] < \pi^{*} < E\left[\tilde{\pi}, \nu^{opt}\left(\cdot \mid \mathbf{I}^{*}\right)\right]$$

and

$$E\left[\tilde{\pi}, \nu^{FB}\left(\cdot \mid \mathbf{I}^{*}\right)\right] \leq \pi^{*} \text{ iff } \lambda \leq \pi^{*}.$$

7 Diverging posteriors and attitude polarization

We are now ready to state and prove our main results whereby we suppose that agents have received the same (limit) sample information from the statistical experiment. To focus our analysis we only consider interesting differences between the agents' learning behavior. In particular, we differentiate between two relevant cases of heterogenous learning behavior. On the one hand, we consider full Bayesian learners who have different initial attitudes with respect to optimism under ambiguity implying different prior beliefs. On the other hand, we consider agents who may have identical prior beliefs but have different, i.e., optimistic resp. pessimistic, attitudes with respect to the interpretation of new information.

Formally, consider a set of agents, I, such that, for every agent $i \in I$, the prior about the parameter π is given by

$$\nu_i\left(\boldsymbol{\pi}\right) = \delta_i \lambda_i + \left(1 - \delta_i\right) \cdot K_{\alpha_i, \beta_i} \pi^{\alpha_i - 1} \left(1 - \pi\right)^{\beta_i - 1} \text{ for } \pi \in [0,1].$$

For the sake of expositional clarity, we restrict attention to the case in which differences in initial beliefs of agents can only be due to their respective optimism parameters λ_i, $i \in I$, under ambiguity.

Assumption 1. *The priors of all agents $i \in I$ satisfy $\delta_i = \delta$, $\alpha_i = \alpha$, and $\beta_i = \beta$ for some parameter values δ, α, β.*

By the following assumption we restrict attention to the interesting case of non-degenerate objective probabilities.

Assumption 2. *The "true" probability π^* is non-degenerate, i.e., $\pi^* \in (0, 1)$.*

While an extension to the case $\pi^* \in [0, 1]$ is straightforward, we avoid by the above assumption the discussion of tedious boundary conditions which would not add to the understanding of our general findings.

As our first main result (proposition 1) we identify conditions under which posterior beliefs diverge such that the directed distance between the posterior beliefs of the two agents is strictly greater than the directed distance between their priors. That is, our first result refers to *diverging posteriors* in the following sense.

Definition (Diverging Posteriors). *Let $I = \{1, 2\}$. We say that both agents'* posteriors strictly diverge *iff*

$$E[\tilde{\pi}, \nu_1(\cdot \mid \mathbf{I}^*)] - E[\tilde{\pi}, \nu_2(\cdot \mid \mathbf{I}^*)] > E[\tilde{\pi}, \nu_1(\cdot)] - E[\tilde{\pi}, \nu_2(\cdot)] \tag{18}$$

whereby

$$E[\tilde{\pi}, \nu_1(\cdot)] \geq E[\tilde{\pi}, \nu_2(\cdot)]. \tag{19}$$

According to our concept of strictly diverging posteriors, the repeated learning of identical information will widen any initial gap in prior beliefs whereby the posteriors may move in the same direction. We also refer to this divergence in beliefs as a *weak* form of myside bias.

Proposition 1. (Diverging Posteriors)

Let $I = \{1, 2\}$ and suppose that assumptions 1 and 2 are satisfied.

(i) *Assume that both agents are full Bayesian learners. Then the agents' posteriors strictly diverge if and only if $\delta > 0$ and $\lambda_1 > \lambda_2$.*

(ii) *Assume that agent* 1 *is an optimistic whereas agent* 2 *is a pessimistic Bayesian learner. Then the agents' posteriors strictly diverge if and only if* $\delta > 0$ *and* $\lambda_1 \geq \lambda_2$.

Observation 6 is useful for understanding both results of the above proposition. The first result is driven by the initial gap of beliefs. The intensification of ambiguity attitudes in the course of the learning process increases this initial gap. The second result is due to the assumption of different learning rules of the agents. While the two agents may have the same prior beliefs, they interpret the same information differentially which then leads to a gap in beliefs in the course of the learning process.

Our second main result (proposition 2) focuses on conditions that ensure attitude polarization. Attitude polarization in our sense is a stronger concept than mere divergence of posteriors in that it additionally requires that the posteriors move in opposite directions. We also refer to this divergence in beliefs as a *strong* form of myside bias. Formally, we consider the following definition of *attitude polarization.*

Definition (Attitude Polarization). *Let* $I = \{1, 2\}$. *We say that both agents'* attitudes become strictly polarized *iff*

$$E\left[\tilde{\pi}, \nu_1\left(\cdot \mid \mathbf{I}^*\right)\right] > E\left[\tilde{\pi}, \nu_1\left(\cdot\right)\right] \geq E\left[\tilde{\pi}, \nu_2\left(\cdot\right)\right] > E\left[\tilde{\pi}, \nu_2\left(\cdot \mid \mathbf{I}^*\right)\right]. \tag{20}$$

In order to further focus our analysis we restrict attention to the case in which the additive part of the prior estimate coincides with the objective probability.

Assumption 3. *The priors of all agents* $i \in I$ *satisfy* $E\left[\tilde{\pi}, \mu\left(\cdot\right)\right] = \pi^*$.

Proposition 2. (Attitude Polarization I)

Let $I = \{1, 2\}$ *and suppose that assumptions 1, 2, and 3 are satisfied.*

(i) *Assume that both agents are full Bayesian learners. Then the agents' attitudes become strictly polarized if and only if* $\delta \in (0, 1)$, $\lambda_1 > \lambda_2$, *and*

$$\lambda_1 > \pi^* > \lambda_2. \tag{21}$$

(ii) *Assume that agent* 1 *is an optimistic whereas agent* 2 *is a pessimistic Bayesian learner. Then the agents' attitudes become strictly polarized if and only if* $\delta > 0$ *and* $\lambda_1 \geq \lambda_2$.

While the second result in the above proposition basically states the same conditions as the result of proposition 1(ii) (under the additional assumption 3), the first result deserves commenting on. Condition (21) has a straightforward interpretation: Whenever one of the two agents resolves his ambiguity in a more optimistic way and the other agent resolves it in a more pessimistic way than the weight of the true objective probability suggests, then the intensification of the initial gap of prior beliefs in the course of the learning process moves the posterior beliefs into opposite directions.

Our formal definitions of "diverging posteriors" and "attitude polarization" capture the idea that the agents' posteriors diverge rather than converge despite the fact that they receive the same information. The results of propositions 1 and 2 demonstrate that this weak, respectively strong, form of a myside bias may occur in different learning scenarios. While the results of propositions 1(i) and 2(i) are driven by the initial gap in prior beliefs, the results of propositions 1(ii) and 2(ii) build upon the different learning rules of the agents. According to condition (21) attitude polarization for full Bayesian learners rather occurs if the difference in initial beliefs is large, i.e., strong optimism of agent 1 versus strong pessimism of agent 2. Such a difference in prior beliefs is not necessary for attitude polarization in case the agents apply different learning rules. That is, even agents with common priors may experience diverging posteriors and attitude polarization if they interpret new information differently.

Finally, the following proposition shows that whenever full Bayesian learners express attitude polarization, the magnitude of attitude polarization between an optimistic and a pessimistic learner will be even more significant. This (intuitive) result is an immediate consequence of observation 7.

Proposition 3. (Attitude Polarization II)

Let $I = \{1, ..., 4\}$ *and suppose that assumptions 1, 2, and 3 are satisfied whereby we have for the agents' priors*

$$\lambda_3 = \lambda_1 > \lambda_2 = \lambda_4.$$

Further assume that agents 1 *and* 2 *are full Bayesian learners whereas agent* 3 *is an optimistic and agent* 4 *is a pessimistic Bayesian learner. If the attitudes of agents*

1 *and* 2 *become strictly polarized, then the attitudes of agents* 3 *and* 4 *are even more polarized, i.e.,*

$$E[\tilde{\pi}, \nu_3(\cdot \mid \mathbf{I}^*)] > E[\tilde{\pi}, \nu_1(\cdot \mid \mathbf{I}^*)] > E[\tilde{\pi}, \nu_2(\cdot \mid \mathbf{I}^*)] > E[\tilde{\pi}, \nu_4(\cdot \mid \mathbf{I}^*)].$$

8 Conclusion

To account for the empirical phenomena of "myside bias" and "irrational belief persistence" in people's learning behavior we propose formal models of Bayesian learning where the interpretation of new information is prone to psychological biases. Based on a simplified representation of ambiguous beliefs we develop parsimonious representations of the agent's initial beliefs and updating processes. Our parsimonious representation is particularly attractive because the deviations from rational Bayesian learning implied by the parametrization of beliefs come along with direct psychological interpretations: One parameter measures the lack of confidence (ambiguity) of the decision maker in his additive prior belief, the second parameter measures the degree of optimism, respectively pessimism, that the decision maker attaches to a resolution of ambiguity in the course of the learning process. We then focus attention on three alternative updating rules that are characterized by different degrees of optimism, respectively pessimism, in the interpretation of new information. As a specific feature of our approach, the resulting models of Bayesian learning with psychological attitudes reduce to a standard model of rational Bayesian learning in the absence of ambiguity. However, we show that this standard model of rational Bayesian learning alone results in convergent beliefs and is therefore not a suitable framework to account for phenomena such as a myside bias. Our model thereby enables us to capture such phenomena within an axiomatically founded decision theoretic framework without resorting to *ad hoc* specifications of psychological biases as, e.g., in Rabin and Schrag (1999).

We then develop a heterogeneous agents setting to derive divergent posterior beliefs and attitude polarization for the agents' learning processes under ambiguity. Attitude polarization is defined as a stronger condition than divergent beliefs in that the posterior beliefs of two agents move into opposite directions. While we assume that the agents receive the same information, the agents may have different prior beliefs or apply different learning rules. Two main findings emerge:

1. We may observe divergent posterior beliefs and attitude polarization for agents who have identical attitudes with respect to the interpretation of new information but have different initial attitudes with respect to optimism, resp. pessimism, under ambiguity.

2. We may observe divergent posterior beliefs and attitude polarization in case the agents have identical initial attitudes with respect to optimism, resp. pessimism, under ambiguity but have different attitudes with respect to the interpretation of new information.

Our stylized models of Bayesian learning with non-additive beliefs thus formally accommodate two alternative scenarios of a "myside bias". In a first scenario, a "myside bias" arises because of personal attitudes towards the resolution of ambiguity. In a second scenario, a "myside bias" corresponds to personal attitudes towards the interpretation of information. While the psychological studies quoted in the introduction provide empirical evidence for the phenomenon of attitude polarization, they do not differentiate between these two alternative explanations for the phenomenon. It would therefore be interesting to gather more empirical evidence on updating and learning with non-additive beliefs. In this respect, our formal model may be useful for designing experiments that specifically look at the issue of Bayesian updating of ambiguous beliefs.

In future research we aim to apply our approach to topics in information economics that are typically analyzed under the assumption of rational Bayesian learning such as fictitious play in strategic games (see, e.g., Fudenberg and Kreps 1993; Fudenberg and Levine 1995; Krishna and Sjostrom 1998) or no-trade results (see, e.g., Milgrom and Stokey 1982; Morris 1994; Neeman 1996; Zimper 2008). Along the line of heterogeneous agent models that depart from the rational expectations or rational Bayesian learning paradigms, our approach may also have promising implications for asset pricing models (see, e.g., Cecchetti, Lam, and Mark 2000; Abel 2002; Ludwig and Zimper 2008b) and theories of endogenous speculative bubbles (see, e.g., the discussion in Kurz 1996).

Appendix: Formal proofs

Proof of observation 2: By an argument in Schmeidler (1986, p. 256f), it suffices to restrict attention to a non-negative valued random variable Y so that

$$E[Y,\nu] = \int_0^{+\infty} \nu\{\omega \in \Omega \mid Y(\omega) \geq z\}\, dz,$$

which is equivalent to

$$E[Y,\nu] = \int_{\min Y}^{\max Y} \nu\{\omega \in \Omega \mid Y(\omega) \geq z\}\, dz \tag{22}$$

since Y is closed and bounded. We consider a partition P_n, $n = 1, 2, \ldots$, of Ω with members

$$A_n^k = \{\omega \in \Omega \mid a_{k,n} < X(\omega) \leq b_{k,n}\} \text{ for } k = 1, \ldots, 2^n$$

such that

$$\begin{aligned} a_{k,n} &= [\max Y - \min Y] \cdot \frac{(k-1)}{2^n} + \min Y \\ b_{k,n} &= [\max Y - \min Y] \cdot \frac{k}{2^n} + \min Y. \end{aligned}$$

Define the step functions $a_n : \Omega \to \mathbb{R}$ and $b_n : \Omega \to \mathbb{R}$ such that, for $\omega \in A_n^k$, $k = 1, \ldots, 2^n$,

$$\begin{aligned} a_n(\omega) &= a_{k,n} \\ b_n(\omega) &= b_{k,n}. \end{aligned}$$

Obviously,

$$E[a_n,\nu] \leq E[Y,\nu] \leq E[b_n,\nu]$$

for all n and

$$\lim_{n\to\infty} E[b_n,\nu] - E[a_n,\nu] = 0.$$

That is, $E[a_n,\nu]$ and $E[b_n,\nu]$ converge to $E[Y,\nu]$ for $n \to \infty$. Furthermore, observe that

$$\begin{aligned} \min a_n &= \min Y \text{ for all } n, \text{ and} \\ \max b_n &= \max Y \text{ for all } n. \end{aligned}$$

Since $\lim_{n\to\infty} \min b_n = \lim_{n\to\infty} \min a_n$ and $E[b_n,\mu]$ is continuous in n, we have

$$\begin{aligned} \lim_{n\to\infty} E[b_n,\nu] &= \delta\left(\lambda \lim_{n\to\infty} \max b_n + (1-\lambda) \lim_{n\to\infty} \min b_n\right) + (1-\delta) \lim_{n\to\infty} E[b_n,\mu] \\ &= \delta(\lambda \max Y + (1-\lambda)\min Y) + (1-\delta) E[Y,\mu]. \end{aligned}$$

In order to prove proposition 3, it therefore remains to be shown that, for all n,

$$E[b_n,\nu] = \delta(\lambda \max b_n + (1-\lambda)\min b_n) + (1-\delta) E[b_n,\mu].$$

Since b_n is a step function, (22) becomes

$$\begin{aligned} E[b_n,\nu] &= \sum_{A_n^k \in P_n} \nu\left(A_n^{2^n} \cup ... \cup A_n^k\right) \cdot (b_{k,n} - b_{k-1,n}) \\ &= \sum_{A_n^k \in P_n} b_{k,n} \cdot \left[\nu\left(A_n^{2^n} \cup ... \cup A_n^k\right) - \nu\left(A_n^{2^n} \cup ... \cup A_n^{k-1}\right)\right], \end{aligned}$$

implying for a neo-additive capacity

$$\begin{aligned} E[b_n,\nu] &= \max b_n \left[\delta\lambda + (1-\delta)\mu\left(A_n^{2^n}\right)\right] + \sum_{k=2}^{2^n-1} b_{k,n}(1-\delta)\mu\left(A_n^k\right) \\ &\quad + \min b_n \left[1 - \delta\lambda - (1-\delta)\sum_{k=2}^{2^n}\mu\left(A_n^k\right)\right] \\ &= \delta\lambda \max b_n + (1-\delta)\sum_{k=1}^{2^n} b_{k,n}\mu\left(A_n^k\right) + \min b_n[\delta - \delta\lambda] \\ &= \delta(\lambda \max b_n + (1-\lambda)\min b_n) + (1-\delta) E[b_n,\mu]. \end{aligned}$$

□

Proof of observation 3: An application of the full Bayesian update rule to a neo-additive capacity gives

$$\begin{aligned} \nu^{FB}(A \mid B) &= \frac{\delta\cdot\lambda + (1-\delta)\cdot\mu(A\cap B)}{\delta\cdot\lambda + (1-\delta)\cdot\mu(A\cap B) + 1 - (\delta\cdot\lambda + (1-\delta)\cdot\mu(A\cup\neg B))} \\ &= \frac{\delta\cdot\lambda + (1-\delta)\cdot\mu(A\cap B)}{1 + (1-\delta)\cdot(\mu(A\cap B) - \mu(A\cup\neg B))} \\ &= \frac{\delta\cdot\lambda + (1-\delta)\cdot\mu(A\cap B)}{1 + (1-\delta)\cdot(\mu(A\cap B) - \mu(A) - \mu(\neg B) + \mu(A\cap\neg B))} \\ &= \frac{\delta\cdot\lambda + (1-\delta)\cdot\mu(A\cap B)}{1 + (1-\delta)\cdot(-\mu(\neg B))} \\ &= \frac{\delta\cdot\lambda + (1-\delta)\cdot\mu(A\cap B)}{\delta + (1-\delta)\cdot\mu(B)} \\ &= \frac{\delta\cdot\lambda}{\delta + (1-\delta)\cdot\mu(B)} + \frac{(1-\delta)\cdot\mu(B)}{\delta + (1-\delta)\cdot\mu(B)}\mu(A\mid B) \\ &= \delta_B^{FB}\cdot\lambda + \left(1-\delta_B^{FB}\right)\cdot\mu(A\mid B) \end{aligned}$$

with δ_B^{FB} given by (10).□

Proof of observation 4: An application of the optimistic Bayesian update rule to a neo-additive capacity gives

$$\begin{aligned}
\nu^{opt}(A \mid B) &= \frac{\delta \cdot \lambda + (1-\delta) \cdot \mu(A \cap B)}{\delta \cdot \lambda + (1-\delta) \cdot \mu(B)} \\
&= \frac{\delta \cdot \lambda}{\delta \cdot \lambda + (1-\delta) \cdot \mu(B)} + \frac{(1-\delta) \cdot \mu(B)}{\delta \cdot \lambda + (1-\delta) \cdot \mu(B)} \cdot \mu(A \mid B) \\
&= \delta_B^{opt} + \left(1 - \delta_B^{opt}\right) \cdot \mu(A \mid B)
\end{aligned}$$

whereby

$$\delta_B^{opt} = \frac{\delta \cdot \lambda}{\delta \cdot \lambda + (1-\delta) \cdot \mu(B)}.$$

□

Proof of observation 5: An application of the pessimistic Bayesian update rule to a neo-additive capacity gives

$$\begin{aligned}
\nu^{pess}(A \mid B) &= \frac{\nu(A \cup \neg B) - \nu(\neg B)}{1 - \nu(\neg B)} \\
&= \frac{\delta \cdot \lambda + (1-\delta) \cdot \mu(A \cup \neg B) - \delta \cdot \lambda - (1-\delta) \cdot \mu(\neg B)}{1 - \delta \cdot \lambda - (1-\delta) \cdot \mu(\neg B)} \\
&= \frac{(1-\delta) \cdot \mu(\neg(\neg A \cap B)) - (1-\delta) \cdot \mu(\neg B)}{1 - \delta \cdot \lambda - (1-\delta) \cdot \mu(\neg B)} \\
&= \frac{(1-\delta) \cdot (1 - \mu(\neg A \cap B)) - (1-\delta) \cdot (1 - \mu(B))}{1 - \delta \cdot \lambda - (1-\delta) \cdot \mu(\neg B)} \\
&= \frac{(1-\delta) \cdot (\mu(B) - \mu(\neg A \cap B))}{1 - \delta \cdot \lambda - (1-\delta) \cdot (\mu(\neg B))} \\
&= \frac{(1-\delta) \cdot (\mu(B) - \mu(B)\, \mu(\neg A \mid B))}{1 - \delta \cdot \lambda - (1-\delta) \cdot (\mu(\neg B))} \\
&= \frac{(1-\delta) \cdot (\mu(B) - \mu(B)\,(1 - \mu(A \mid B)))}{1 - \delta \cdot \lambda - (1-\delta) \cdot \mu(\neg B)} \\
&= (1 - \delta_B^{pess}) \cdot \mu(A \mid B)
\end{aligned}$$

whereby

$$\delta_B^{pess} = \frac{\delta(1-\lambda)}{\delta(1-\lambda) + (1-\delta) \cdot \mu(B)}.$$

□

Proof of observation 6: This observation is an immediate consequence of the following lemma combined with the fact that the additive part of the neo-additive beliefs converges in probability to the true probability π^*, i.e.,

$$\lim_{n\to\infty} prob\left(\left|E\left[\tilde{\pi}, \mu\left(\cdot \mid \mathbf{I}_n^k\right)\right] - \pi^*\right| \leq c\right) = 1$$

for some $c > 0$. The lemma simply uses our results (observations 3-5) on Bayesian updating of neo-additive capacities in order to derive conditional neo-additive capacities for the special case (17).

Lemma. *Suppose that the agent receives sample information $\mathbf{I}_n^k$. Contingent on the applied update rule we obtain the following conditional neo-additive beliefs and posterior estimates about parameter π whereby $E\left[\tilde{\pi}, \mu\left(\cdot \mid \mathbf{I}_n^k\right)\right]$ is given by (1).*

(i) *Full Bayesian updating.*

$$\nu^{FB}\left(\boldsymbol{\pi} \mid \mathbf{I}_n^k\right) = \delta_{\mathbf{I}_n^k}^{FB}\lambda + \left(1 - \delta_{\mathbf{I}_n^k}^{FB}\right) \cdot K_{\alpha+k,\beta+n-k}\pi^{\alpha+k-1}\left(1-\pi\right)^{\beta+n-k-1}$$

with

$$\delta_{\mathbf{I}_n^k}^{FB} = \frac{\delta}{\delta + (1-\delta) \cdot \mu\left(\mathbf{I}_n^k\right)}$$

so that

$$E\left[\tilde{\pi}, \nu^{FB}\left(\cdot \mid \mathbf{I}_n^k\right)\right] = \delta_{\mathbf{I}_n^k}^{FB}\lambda + \left(1 - \delta_{\mathbf{I}_n^k}^{FB}\right) \cdot E\left[\tilde{\pi}, \mu\left(\cdot \mid \mathbf{I}_n^k\right)\right].$$

(ii) *Optimistic Bayesian updating.*

$$\nu^{opt}\left(\boldsymbol{\pi} \mid \mathbf{I}_n^k\right) = \delta_{\mathbf{I}_n^k}^{opt} + \left(1 - \delta_{\mathbf{I}_n^k}^{opt}\right) \cdot K_{\alpha+k,\beta+n-k}\pi^{\alpha+k-1}\left(1-\pi\right)^{\beta+n-k-1}$$

with

$$\delta_{\mathbf{I}_n^k}^{opt} = \frac{\delta \cdot \lambda}{\delta \cdot \lambda + (1-\delta) \cdot \mu\left(\mathbf{I}_n^k\right)}$$

so that

$$E\left[\tilde{\pi}, \nu^{opt}\left(\cdot \mid \mathbf{I}_n^k\right)\right] = \delta_{\mathbf{I}_n^k}^{opt} + \left(1 - \delta_{\mathbf{I}_n^k}^{opt}\right) \cdot E\left[\tilde{\pi}, \mu\left(\cdot \mid \mathbf{I}_n^k\right)\right].$$

(iii) *Pessimistic Bayesian updating.*

$$\nu^{pess}\left(\boldsymbol{\pi} \mid \mathbf{I}_n^k\right) = \left(1 - \delta_{\mathbf{I}_n^k}^{pess}\right) \cdot K_{\alpha+k,\beta+n-k}\pi^{\alpha+k-1}\left(1-\pi\right)^{\beta+n-k-1}$$

with

$$\delta_{\mathbf{I}_n^k}^{pess} = \frac{\delta \cdot (1-\lambda)}{\delta \cdot (1-\lambda) + (1-\delta) \cdot \mu\left(\mathbf{I}_n^k\right)}$$

so that

$$E\left[\tilde{\pi}, \nu^{pess}\left(\cdot \mid \mathbf{I}_n^k\right)\right] = \left(1 - \delta_{\mathbf{I}_n^k}^{pess}\right) \cdot E\left[\tilde{\pi}, \mu\left(\cdot \mid \mathbf{I}_n^k\right)\right].$$

Proof of observation 7: At first observe that $\delta > 0$ and $\lambda \in (0,1)$ implies $\delta_{\mathbf{I}^*}^{FB} > \delta_{\mathbf{I}^*}^{opt}$ as well as $\delta_{\mathbf{I}^*}^{FB} > \delta_{\mathbf{I}^*}^{pess}$. Consider the inequality

$$\begin{aligned} E\left[\tilde{\pi}, \nu^{FB}\left(\cdot \mid \mathbf{I}^*\right)\right] &< E\left[\tilde{\pi}, \nu^{opt}\left(\cdot \mid \mathbf{I}^*\right)\right] \Leftrightarrow \\ \delta_{\mathbf{I}^*}^{FB} \cdot \lambda &< \delta_{\mathbf{I}^*}^{opt} + \left(\delta_{\mathbf{I}^*}^{FB} - \delta_{\mathbf{I}^*}^{opt}\right) \cdot \pi^*, \end{aligned}$$

which holds, by $\delta_{\mathbf{I}^*}^{FB} > \delta_{\mathbf{I}^*}^{opt}$, for all π^* iff

$$\begin{aligned} \delta_{\mathbf{I}^*}^{FB} \cdot \lambda &< \delta_{\mathbf{I}^*}^{opt} \Leftrightarrow \\ \frac{\delta \cdot \lambda}{\delta + (1-\delta) \cdot \mu\left(\mathbf{I}^*\right)} &< \frac{\delta \cdot \lambda}{\delta \cdot \lambda + (1-\delta) \cdot \mu\left(\mathbf{I}^*\right)} \Leftrightarrow \\ \lambda &< 1. \end{aligned}$$

Turn now to the inequality

$$\begin{aligned} E\left[\tilde{\pi}, \nu^{pess}\left(\cdot \mid \mathbf{I}^*\right)\right] &< E\left[\tilde{\pi}, \nu^{FB}\left(\cdot \mid \mathbf{I}^*\right)\right] \Leftrightarrow \\ \left(\delta_{\mathbf{I}^*}^{FB} - \delta_{\mathbf{I}^*}^{pess}\right) \cdot \pi^* &< \delta_{\mathbf{I}^*}^{FB} \cdot \lambda, \end{aligned}$$

which holds, by $\delta_{\mathbf{I}^*}^{FB} > \delta_{\mathbf{I}^*}^{pess}$, for all π^* iff

$$\begin{aligned} \left(\delta_{\mathbf{I}^*}^{FB} - \delta_{\mathbf{I}^*}^{pess}\right) &< \delta_{\mathbf{I}^*}^{FB} \cdot \lambda \Leftrightarrow \\ \delta_{\mathbf{I}^*}^{FB} \cdot (1-\lambda) &< \delta_{\mathbf{I}^*}^{pess} \Leftrightarrow \\ \frac{\delta \cdot (1-\lambda)}{\delta + (1-\delta) \cdot \mu\left(\mathbf{I}^*\right)} &< \frac{\delta \cdot (1-\lambda)}{\delta \cdot (1-\lambda) + (1-\delta) \cdot \mu\left(\mathbf{I}^*\right)} \Leftrightarrow \\ 0 &< \lambda. \end{aligned}$$

This proves the first part of the observation. The second part readily follows from the assumption that $\mu\left(\mathbf{I}^*\right) < 1.\square$

Proof of proposition 1.

Part (i). Observe at first that inequality (19) is satisfied if and only if $\lambda_1 \geq \lambda_2$. Obviously, if $\lambda_1 = \lambda_2$ then (18) must be violated. Thus we can restrict attention to $\lambda_1 > \lambda_2$. Observe that, by the corollary, (18) writes as

$$\begin{aligned} & \delta_{\mathbf{I}^*}^{FB} \lambda_1 + \left(1 - \delta_{\mathbf{I}^*}^{FB}\right) \cdot \pi^* - \delta_{\mathbf{I}^*}^{FB} \lambda_2 + \left(1 - \delta_{\mathbf{I}^*}^{FB}\right) \cdot \pi^* \\ > \; & \delta \cdot \lambda_1 + (1-\delta) \cdot E\left[\tilde{\pi}, \mu(\cdot)\right] - \left(\delta \cdot \lambda_2 + (1-\delta) \cdot E\left[\tilde{\pi}, \mu(\cdot)\right]\right), \end{aligned}$$

which is equivalent to

$$\begin{aligned} \delta_{\mathbf{I}^*}^{FB} &> \delta \Leftrightarrow \\ \frac{\delta}{\delta + (1-\delta) \cdot \mu\left(\mathbf{I}^*\right)} &> \delta, \end{aligned}$$

and therefore holds if and only if $\delta \in (0,1)$ since $\mu(\mathbf{I}^*) < 1.\square$

Part (ii). Again, observe at first that inequality (19) is satisfied if and only if $\lambda_1 \geq \lambda_2$. By the corollary, (18) becomes

$$\begin{aligned} & \delta_{\mathbf{I}^*}^{opt} + \left(1 - \delta_{\mathbf{I}^*}^{opt}\right) \cdot \pi^* - \left(1 - \delta_{\mathbf{I}^*}^{pess}\right) \cdot \pi^* \\ > \; & \delta \cdot \lambda_1 + (1-\delta) \cdot E\left[\tilde{\pi}, \mu(\cdot)\right] - \left(\delta \cdot \lambda_2 + (1-\delta) \cdot E\left[\tilde{\pi}, \mu(\cdot)\right]\right) \end{aligned}$$

which is equivalent to

$$\delta_{\mathbf{I}^*}^{opt} + \left(\delta_{\mathbf{I}^*}^{pess} - \delta_{\mathbf{I}^*}^{opt}\right) \cdot \pi^* > \delta\left(\lambda_1 - \lambda_2\right). \tag{23}$$

If $\delta = 0$, the l.h.s. as well as the r.h.s. of (23) equal zero. Thus, $\delta > 0$ is a necessary condition for (23) to hold. In what follows we prove that $\delta > 0$ is also a sufficient condition. Let $\delta > 0$ and consider at first the case that

$$\delta_{\mathbf{I}^*}^{pess} - \delta_{\mathbf{I}^*}^{opt} \leq 0. \tag{24}$$

Since the l.h.s. of (23) is then continuously strictly decreasing in π^* and, by assumption, $\pi^* \in (0,1)$, (23) is satisfied for all π^* if and only if

$$\begin{gathered} \delta_{\mathbf{I}^*}^{pess} \geq \delta\left(\lambda_1 - \lambda_2\right) \Leftrightarrow \\ \frac{1-\lambda_2}{\delta(1-\lambda_2) + (1-\delta)\mu(\mathbf{I}^*)} \geq \lambda_1 - \lambda_2, \end{gathered}$$

which is obviously true for all λ_1, λ_2 since

$$\frac{1-\lambda_2}{\delta(1-\lambda_2) + (1-\delta)\mu(\mathbf{I}^*)} \geq 1 - \lambda_2.$$

This proves the claim for case (24).

Let $\delta > 0$ and consider now the converse case

$$\delta_{\mathbf{I}^*}^{pess} - \delta_{\mathbf{I}^*}^{opt} > 0. \tag{25}$$

Since the l.h.s. of (23) is then continuously strictly increasing in $\pi^* \in (0,1)$, (23) is satisfied for all π^* if and only if

$$\begin{aligned} \delta_{\mathbf{I}^*}^{opt} + \left(\delta_{\mathbf{I}^*}^{pess} - \delta_{\mathbf{I}^*}^{opt}\right) & \geq \delta\left(\lambda_1 - \lambda_2\right) \Leftrightarrow \\ \delta_{\mathbf{I}^*}^{pess} & \geq \delta\left(\lambda_1 - \lambda_2\right) \Leftrightarrow \\ \frac{1-\lambda_2}{\delta(1-\lambda_2) + (1-\delta)\mu(\mathbf{I}^*)} & \geq \lambda_1 - \lambda_2, \end{aligned}$$

which is obviously true for all λ_1, λ_2. This proves that $\delta > 0$ is sufficient for (23) to hold. $\square\square$

Proof of proposition 2.

Part (i). By the corollary and the assumption that $E[\tilde{\pi}, \mu(\cdot)] = \pi^*$ equation (20) implies

$$E[\tilde{\pi}, \nu_1(\cdot \mid \mathbf{I}^*)] > E[\tilde{\pi}, \nu_2(\cdot \mid \mathbf{I}^*)] \Leftrightarrow$$
$$\delta_{\mathbf{I}^*}^{FB}\lambda_1 + \left(1 - \delta_{\mathbf{I}^*}^{FB}\right) \cdot \pi^* > \delta_{\mathbf{I}^*}^{FB}\lambda_2 + \left(1 - \delta_{\mathbf{I}^*}^{FB}\right) \cdot \pi^*$$

which holds if and only if $\lambda_1 > \lambda_2$ so that the middle inequality in (20) is also strict. Focus now on the inequalities

$$E[\tilde{\pi}, \nu_1(\cdot \mid \mathbf{I}^*)] > E[\tilde{\pi}, \nu_1(\cdot)] \Leftrightarrow$$
$$\delta_{\mathbf{I}^*}^{FB}\lambda_1 + \left(1 - \delta_{\mathbf{I}^*}^{FB}\right) \cdot \pi^* > \delta \cdot \lambda_1 + (1 - \delta) \cdot \pi^*$$

and

$$E[\tilde{\pi}, \nu_2(\cdot)] > E[\tilde{\pi}, \nu_2(\cdot \mid \mathbf{I}^*)] \Leftrightarrow$$
$$\delta \cdot \lambda_2 + (1 - \delta) \cdot \pi^* > \delta_{\mathbf{I}^*}^{FB}\lambda_2 + \left(1 - \delta_{\mathbf{I}^*}^{FB}\right) \cdot \pi^*$$

which are implied by (20) under the assumption that $E[\tilde{\pi}, \mu(\cdot)] = \pi^*$. Observe that these inequalities require $\delta \in (0, 1)$ since $\delta \in \{0, 1\}$ would imply $\delta_{\mathbf{I}^*}^{FB} = \delta$. As a consequence of $\delta \in (0, 1)$, we have from the corollary that $\delta_{\mathbf{I}^*}^{FB} > \delta$ because $\mu(\mathbf{I}^*) < 1$ so that the above inequalities hold if and only if

$$\lambda_1 > \pi^* > \lambda_2,$$

which proves the result. □

Part (ii). By the corollary, the inequality $E[\tilde{\pi}, \nu_1(\cdot)] \geq E[\tilde{\pi}, \nu_2(\cdot)]$ in (20) holds if and only if $\lambda_1 \geq \lambda_2$. Consider at first agent 1 and rewrite the relevant part in (20) as

$$\begin{aligned} E\left[\tilde{\pi}, \nu_1^{opt}(\cdot \mid \mathbf{I}^*)\right] &> E[\tilde{\pi}, \nu_1(\cdot)] \Leftrightarrow \\ \delta_{\mathbf{I}^*}^{opt} + \left(1 - \delta_{\mathbf{I}^*}^{opt}\right) \cdot \pi^* &> \delta \cdot \lambda_1 + (1 - \delta) \cdot E[\tilde{\pi}, \mu(\cdot)] \end{aligned}$$

which, under the assumption that $E[\tilde{\pi}, \mu(\cdot)] = \pi^*$ is equivalent to

$$\delta_{\mathbf{I}^*}^{opt} + \left(\delta - \delta_{\mathbf{I}^*}^{opt}\right) \cdot \pi^* > \delta \cdot \lambda_1. \tag{26}$$

Observe that $\delta > 0$ is a necessary condition for (26) to hold. In what follows we prove that $\delta > 0$ is also sufficient. Let $\delta > 0$ and consider at first the case that

$$\delta - \delta_{\mathbf{I}^*}^{opt} \leq 0. \tag{27}$$

Since the l.h.s. of (26) is then continuously strictly decreasing in $\pi^* \in (0,1)$, (26) is satisfied for all π^* if and only if

$$\delta_{\mathbf{I}^*}^{opt} \geq \delta \cdot \lambda_1 \Leftrightarrow$$
$$\frac{1}{\delta\lambda_1 + (1-\delta)\,\mu(\mathbf{I}^*)} \geq 1,$$

which is obviously satisfied. This proves the claim for case (27).

Let $\delta > 0$ and consider now the converse case

$$\delta - \delta_{\mathbf{I}^*}^{opt} > 0. \tag{28}$$

Since the l.h.s. of (26) is then continuously strictly increasing in $\pi^* \in (0,1)$, (26) is satisfied for all π^* if and only if

$$\delta \geq \delta \cdot \lambda_1,$$

which is obviously satisfied. This proves our claim that $\delta > 0$ is a necessary and sufficient condition for (26) to hold.□□

References

Abel, A.B. (2002), "An Exploration of the Effects of Pessimism and Doubt on Asset Returns", *Journal of Economic Dynamics and Control* **26**, 1075-1092.

Acemoglu, D., V. Chernozhukov and M. Yildiz (2007), "Learning and Disagreement in an Uncertain World ", Working Paper, MIT.

Alesina, A., Baqir, R. and Easterly, W. (1999), "Public Goods and Ethnic Divisions", *Quarterly Journal of Economics* **115(4)**, 1167-1199.

Baron, J. (2008), *Thinking and Deciding*, Cambridge University Press: New York, Melbourne, Madrid.

Blackwell, D. and L. Dubins (1962), "Merging of Opinions with Increasing Information", *The Annals of Mathematical Statistics* **33**, 882-886.

Brav, A. and J.B. Heaton (2002), "Competing Theories of Financial Anomalies", *Review of Financial Studies* **15**, 575-606.

Breiman, D., LeCam, L. and L. Schwartz (1964), "Consistent Estimates and Zero-One Sets", *The Annals of Mathematical Statistics* **35**, 157-161.

Cecchetti, S.G, Lam, P. and N.C. Mark (2000), "Asset Pricing with Distorted Beliefs: Are Equity Returns Too Good to Be True?", *American Economic Review* **90**, 787-805.

Chapman, C.R. (1973), "Prior Probability Bias in Information Seeking and Opinion Revision," *The American Journal of Psychology* **86**, 269-282.

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

Dixit, A.K. and J.W. Weibull (2007), "Political Polarization", *Proceedings of the National Academy of Sciences of the United States of America*, published online, doi:10.1073/pnas.0702071104.

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 In Le Calcul des Probabilite´s et ses Applications, Colloques Internationaux du Centre National de la Recherche Scientifique **13**, 23-27.

Duclos, J-Y., J. Esteban and D. Ray (2004), "Polarization: Concepts, Measurement, Estimation", *Econometrica* **72(6)**, 1737-1772.

Eichberger, J. and D. Kelsey (1999), "E-Capacities and the Ellsberg Paradox", *Theory and Decision* **46**, 107-140.

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. (1999), "A Definition of Uncertainty Aversion", *The Review of Economic Studies* **66**, 579-608.

Epstein, L.G. and T. Wang (1994), "Intertemporal Asset Pricing Under Knightian Uncertainty", *Econometrica* **62**, 283-322.

Epstein, L.G. and M. Schneider (2003), "Recursive Multiple-Priors ", *Journal of Economic Theory* **113**, 1-31.

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

Esteban, J. and D. Ray (1994), "On the Measurement of Polarization", *Econometrica*, **62**, 819-852.

Esteban, J. and D. Ray (2001), "Conflict and Distribution", *Journal of Economic Theory*, **87**, 379-415.

Fernández, R. and G. Levy (2008), "Diversity and Redistribution", *Journal of Public Economics*, **92(6)**, 925-943.

Freedman, D.A. (1963), "On the Asymptotic Behavior of Bayes' Estimates in the Discrete Case", *The Annals of Mathematical Statistics* **33**, 882-886.

Fudenberg, D. and D. Kreps (1993), "Learning Mixed Equilibria", *Games and Economic Behavior* **5**, 320-367.

Fudenberg, D. and D.K. Levine (1995), "Consistency and Cautious Fictitious Play", *Journal of Economic Dynamics and Control* **19**, 1065-1090.

Ghirardato, P. and M. Marinacci (2002), "Ambiguity Made Precise: A Comparative Foundation", *Journal of Economic Theory* **102**, 251-289.

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

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

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

Kandel, E. and N.D. Pearson (1995), "Differential Interpretation of Public Signals and Trade in Speculative Markets ", *Journal of Political Economy* **103**, 831-872.

Krishna, V. and T. Sjostrom (1998), "On the Convergence of Fictitious Play", *Mathematics of Operations Research* **23**, 479-511.

Kurz, M. (1994a), "On the Structure and Diversity of Rational Beliefs", *Economic Theory* **4**, 877-900.

Kurz, M. (1994b), "On Rational Belief Equilibrias", *Economic Theory* **4**, 859-876.

Kurz, M. (1996), " Rational Beliefs and Endogenous Uncertainty", *Economic Theory* **8**, 383-397.

Lewellen, J. and J. Shanken (2002), "Learning, Asset-Pricing Tests, and Market Efficiency", *Journal of Finance* **57**, 1113-1145.

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

Lindqvist, E and R. Östling (2008), "Political Polarization and the Size of Government", SSE/EFI Working Paper Series in Economics and Finance No. 628.

Lord, C.G., Ross, L. and M.R. Lepper (1979), "Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence", *Journal of Personality and Social Psychology*, **37**, 2098-2109.

Ludwig, A. and A. Zimper (2008a), "A Parsimonious Choquet Model of Subjective Life-Expectancy", Working Paper.

Ludwig, A. and A. Zimper (2008b), "Biased Bayesian Learning and the Risk-free Rate Puzzle", Working Paper.

Marinacci, M. (1999), "Limit Laws for Non-additive Probabilities and Their Frequentist Interpretation," *Journal of Economic Literature* **84**, 145-195.

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

Milgrom, P. and N. Stockey (1982), "Information, Trade and Common Knowledge", *Journal of Economic Theory* **26**, 17-27.

Morris, S. (1994), "Trade with Heterogeneous Prior Beliefs and Asymmetric Information", *Econometrica* **62**, 1327-1347.

Neeman, Z. (1996), "Approximating Agreeing to Disagree Results with Common p-Beliefs", *Games and Economic Behavior* **16**, 77-96.

Pires, C.P. (2002) "A Rule for Updating Ambiguous Beliefs ", *Theory and Decision* **53**, 137-152.

Pitz, G.F. (1969), "An Inertia Effect (Resistance to Change) in the Revision of Opinion", *Canadian Journal of Psychology* **23**, 24-33.

Pitz, G.F., Downing, L. and H. Reinhold (1967), "Sequential Effects in the Revision of Subjective Probabilities", *Canadian Journal of Psychology* **21**, 381-393.

Rabin, M and J.L. Schrag (1999), "First Impressions Matter: A Model of Confirmatory Bias", *Quarterly Journal of Economics*, **114(1)**, 37-82.

Rudin, W. (1976), *Principles of Mathematical Analysis*, McGraw-Hill: New York.

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. (1986), "Integral Representation without Additivity", *Proceedings of the American Mathematical Society* **97**, 255-261.

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

Siniscalchi, M. (2006), "Dynamic Choice under Ambiguity", mimeo.

Tonks, I. (1983), "Bayesian Learning and the Optimal Investment Decision of the Firm", *The Economic Journal* **93**, 87-98.

Viscusi, W.K. (1985), "A Bayesian Perspective on Biases in Risk Perception", *Economics Letters* **17**, 59-62.

Viscusi, W.K. and C.J. O'Connor (1984), "Adaptive Responses to Chemical Labeling: Are Workers Bayesian Decision Makers?", *The American Economic Review* **74**, 942-956.

Wakker, P.P. (2001), "Testing and Characterizing Properties of Nonadditive Measures through Violations of the Sure-Thing Principle", *Econometrica* **69**, 1039-1059.

Wakker, P.P (2004), "On the Composition of Risk Preference and Belief", *Psychological Review* **111**, 236-241.

Wakker, P.P (2009), Prospect Theory for Risk and Ambiguity, Cambridge University Press, forthcoming.

Weitzman, M. (2007), "Subjective Expectations and Asset-Return Puzzles", *American Economic Review*, forthcoming.

Zimper, A. (2008), "Half-Empty, Half-Full and Why We Can Agree to Disagree Forever", *Journal of Economic Behavior and Organization*, forthcoming.