Half empty, half full and the possibility of *agreeing to disagree*

Alexander Zimper[1]

Working Paper Number 58

[1] School of Economics, University of Cape Town

Half empty, half full and the possibility of *agreeing to disagree*[*]

Alexander Zimper[†]

October 31, 2007

Abstract

Aumann (1976) derives his famous *we cannot agree to disagree* result under the assumption of rational Bayesian learning. Motivated by psychological evidence against this assumption, we develop formal models of *optimistically*, resp. *pessimistically*, *biased Bayesian learning* within the framework of Choquet expected utility theory. As a key feature of our approach the posterior subjective beliefs do, in general, not converge to "true" probabilities. Moreover, the posteriors of different people can converge to different beliefs even if these people receive the same information. As our main contribution we show that people may well *agree to disagree* if their Bayesian learning is psychologically biased in our sense. Remarkably, this finding holds regardless of whether people with identical priors apply the same psychologically biased Bayesian learning rule or not. A simple example about the possibility of ex-post trading in a financial asset illustrates our formal findings. Finally, our analysis settles a discussion in the no-trade literature (cf. Dow, Madrigal, and Werlang 1990, Halevy 1998) in that it clarifies that ex-post trade between agents with common priors and identical learning rules is only possible under asymmetric information.

Keywords: Common Knowledge, No-Trade Results, Rational Bayesian Learning, Bounded Rationality, Choquet Expected Utility Theory, Bayesian Updating, Dynamic Inconsistency

JEL Classification Numbers: D81, G20.

[*]I would like to thank Alexander Ludwig and Elias Khalil for helpful comments and suggestions.

[†]School of Economics, Faculty of Commerce, University of Cape Town, Private Bag, Rondebosch, 7701, Cape Town, RSA. Email: zimper@bigfoot.com.

1 Introduction

1.1 Motivation

Aumann (1976) proves that "If two people have the same priors, and their posteriors for an event A are common knowledge, then these posteriors are equal" (p. 1236). This celebrated *we cannot agree to disagree* result has been derived under the assumption that people's posterior beliefs result from rational Bayesian learning. However, several studies in the psychological literature show that real-life agents systematically violate this assumption in that their learning behavior is prone to effects such as "myside bias" or "irrational belief persistence" (cf., e.g., the references in chapter 9, Baron 2007). For example, in an early contribution to this literature, Lord, Ross, and Lepper (1979) conduct an experiment in which agents' posteriors diverge despite the fact that all agents have received the same information.[1] Moreover, definitions of several psychological phenomenons such as delusions, depressions etc. are based on the observation that different subjects may interpret identical information in different ways (cf. Beck 1976).

The contributions of this paper are two-fold. At first, we introduce in this paper a formal model of Bayesian learning with a myside bias. More precisely, we introduce the notion of optimistically, resp. pessimistically, biased Bayesian learning. With these definitions we formally describe the difference between "half empty" versus "half full" attitudes in the context of interpreting new information. In contrast to the standard model of rational Bayesian learning (e.g., Tonks 1983, Viscusi and O'Connor 1984, Viscusi 1985), the posterior beliefs of biased Bayesian learning do not converge to the "true" probabilities. Furthermore, the posterior beliefs of different agents may well converge to different beliefs even if these agents always receive the same information. As our second contribution, we demonstrate that people may agree to disagree if their Bayesian learning is psychologically biased. This result holds regardless of whether people apply the same Bayesian learning rule or not.

Key to our analysis is the assumption that people are Choquet expected utility (CEU) rather than expected utility (EU) decision makers. CEU theory (Schmeidler 1989, Gilboa 1987) is a generalization of EU theory that admits for the integration of a vNM function with respect to non-additive probability measures (capacities). Properties of such capacities are used for the formal description of ambiguity attitudes which may explain Ellsberg (1961) paradoxes. Ellsberg paradoxes demonstrate systematic violations

[1] The subjects in this experiment were confronted with two purported statistical studies, one study supporting the other study rejecting the claim that capital punishment has a crime deterrence effect. For analogous results in the context of Bayesian updating of subjective probabilities see Pitz, Downing, and Reinhold (1967) and Pitz (1969).

of Savage's (1954) "sure thing principle". The sure thing principle, however, ensures that there is a unique way of deriving ex-post preferences from ex-ante preferences, implying a unique Bayesian update rule for the additive probabilities of subjective EU theory. The picture is different for the non-additive probability measures of CEU theory for which several perceivable Bayesian update rules exist (cf. Gilboa and Schmeidler 1993, Eichberger, Grant and Kelsey 2006, Siniscalchi 2001, 2006). Following Gilboa and Schmeidler's (1993) psychological interpretation we consider the extreme cases of the optimistic, resp. pessimistic, update rule, which we apply to non-additive probability measures defined as neo-additive capacities in the sense of Chateauneuf, Eichberger and Grant (2006). Our resulting definition of optimistically, resp. pessimistically, biased agents combines the standard model of rational Bayesian learning with an optimistic, resp. pessimistic, attitude towards the interpretation of new information.

We present two different results of the type that people may agree to disagree if their learning rules are psychologically biased. Our first result (proposition 1) shows that if two people have the same prior, apply different learning rules, and their posteriors for an event $A \notin \{\varnothing, \Omega\}$ are common knowledge, then these posteriors will be different even in case they have identical information partitions. Within the appropriate framework this result is easily derived. However, beyond the mere formal result our finding addresses an important behavioral issue. Aumann (1976) writes "In private conversation, Tversky has suggested that people may also be biased because of psychological factors, that may make them disregard information that is unpleasant or does not conform to previously formed notions" (p. 1238). There is no way of describing such psychological biases of real-life people within Aumann's framework. Within our approach, however, the resulting "myside bias" has a straightforward interpretation as people's different attitudes towards the interpretation of information due to psychological predispositions such as the "half-empty glass" versus the "half-full glass" attitude.

Whereas our first result applies to people who use different rules of Bayesian learning, our second result (proposition 2) refers to the case of identical learning rules. We find that if two people have the same prior, apply the same learning rule, and their posteriors for an event $A \notin \{\varnothing, \Omega\}$ are common knowledge, then these posteriors can be different in case they have different information partitions. Thus, neither in the case where people have the same information partitions nor in the case where people apply the same update rule does Aumann's conclusion obtain when Bayesian learning is psychologically biased in our sense. To the contrary, according to our results a difference in posteriors that are common knowledge is the rule rather than the exception when people are psychologically biased.

1.2 Relationship to no-trade results

Combined with Harsanyi's (1967) common priors doctrine Aumann's *we cannot agree to disagree* result has been very influential in information economics. Especially the so-called no-trade theorems - basically stating that there should be no ex-post trade in financial assets if the agents are rational - are based on Aumann's approach (cf., e.g., Milgrom 1981, Milgrom and Stokey 1982, Samet 1990, Morris 1994, Bonanno and Nehring 1999). The connection between Aumann's *we cannot agree to disagree* result and the impossibility of ex-post trade in financial assets is straightforward. Under the assumption that agents have different preferences for such assets if and only if they have different beliefs about the assets' future returns, there are strict incentives for ex-post trade if and only if the agents have different posterior beliefs. Since the market-price of such assets is common knowledge between the trading agents, any trade would result in the traders' common knowledge that their posteriors must be different.[2]

Since no trade-results are seemingly at odds with reality, there are several contributions in the literature investigating the robustness of no-trade results with respect to a weakening of Aumann's assumptions. One line of research discusses concepts of bounded rationality that weaken the rationality assumptions of Aumann's epistemic framework. For example, information structures have been considered that are non-partitional (Bacharach 1985, Samet 1990, Geneakoplos 1992, Rubinstein and Wolinsky 1990) or concepts of "almost" common-knowledge have been introduced (Neeman 1996). In contrast to this literature our approach fully adopts Aumann's epistemic framework. The agents of our model are boundedly rational not with respect to their logical capability but with respect to their psychological bias in interpreting new information.

Closer to our own approach is a second line of research on no-trade results that considers decision theoretic alternatives to EU theory. In an early contribution Dow, Madrigal and Werlang (1990) already provide an example in which ex-post trading becomes possible because agents update their non-additive beliefs according to the Dempster-Shafer rule which is at the heart of our definition of pessimistically biased Bayesian learning. Dow et al. thereby assume asymmetric information and common non-additive priors so that their example can be regarded as an illustration of our proposition 2 for the special case of pessimistically biased agents.

Halevy (1998, 2004) claims that the finding of Dow et al. can be extended to the case of symmetric information so that there might occur ex-post trading between agents with common priors and identical information partitions if their beliefs are non-additive. More precisely, Halevy writes:

[2] Note that the typical assumption of "strictly risk averse" traders (e.g., Milgrom 1981, Milgrom and Stokey 1982) is thus not necessary for obtaining "no-trade" results.

> "A similar result appears in Dow et al (1990). Their result, as noted by Epstein and Le Breton (1993) and as our present example illustrates, relies merely on dynamic inconsistency of the individual agents. Their claim that trade is a result of asymmetric information is not accurate: we show below that it could be reached with completely symmetric information and even with a common prior." (footnote 17, p. 20 in Halevy 1998)

In the light of our propositions 1 and 2, we take a somewhat different view from Halevy. Namely, Dow et al.'s conclusion is indeed accurate under their assumption of an identical update rule for all agents: our analysis demonstrates that agents with an identical update rule cannot agree to disagree if they have identical information partitions. Our own asset-trade example in section 5 of this paper therefore establishes the existence of ex-post trade between agents with symmetric information and common priors if and only if the agents have different update rules. Since Halevy's example does not consider different update rules, his finding is apparently at odds with our own results. As it turns out the difference between Dow et al. and our conclusion, on the one hand, and Halevy's conclusion, on the other hand, is due to different notions of common priors. Halevy's example is based on Yaari's (1987) dual theory in which additive probabilities are transformed into non-additive beliefs by some transformation function. Halevy speaks of common priors because his agents have common additive probabilities. But since these agents apply different transformation functions, their resulting non-additive priors are no longer identical. According to our approach and the approach of Dow et al. the assumption of common priors is therefore violated in Halevy's example.

Rubinstein and Wolinsky (1990, Remark p. 190) argue that Milgrom and Stokey's no-trade result applies to all decision theories under uncertainty which satisfy dynamic consistency. Halevy (2004) reports the interesting fact that there might even be ex-post trading between dynamically consistent agents if these agents violate consequentialism. While EU decision makers satisfy, by the sure-thing principle, dynamic consistency as well as consequentialism, the agents of our model only satisfy consequentialism and the possibility of agreeing to disagree exclusively results from their dynamically inconsistent preferences (cf. Epstein and Le Breton 1993, Sarin and Wakker 1998).

The subsequent analysis is structured as follows. In section 2 we describe our decision-theoretic framework which we combine in section 3 with the standard model of rational Bayesian learning. Section 4 recalls Aumann's (1976) epistemic framework and presents our first *agreeing to disagree* result. A simple example in section 5 about the possibility of ex-post asset trade illustrates this first result. Our second *agreeing to disagree* result is stated and proved in section 6. Section 7 concludes.

2 Preliminaries: The decision-theoretic framework

As in Aumann (1976) we consider a measurable space $(\Omega, \mathcal{B})$ with $\mathcal{B}$ denoting a σ-algebra on the state space Ω. As a generalization of Aumann's assumption of EU decision makers, however, we consider a CEU rather than an EU decision maker.[3] In contrast to EU theory, CEU theory admits for non-additive probability measures, i.e., capacities, whereby a capacity $\nu : \mathcal{B} \rightarrow [0,1]$ must satisfy

(i) $\nu(\emptyset) = 0$, $\nu(\Omega) = 1$

(ii) $A \subset B \Rightarrow \nu(A) \leq \nu(B)$ for all $A, B \in \mathcal{B}$.

Additional properties of capacities are used in the literature for formal definitions of, e.g., *ambiguity* and *uncertainty attitudes* (Schmeidler, 1989; Epstein, 1999; Ghirardato and Marinacchi, 2002), *pessimism* and *optimism* (Eichberger and Kelsey, 1999; Wakker, 2001), as well as *sensitivity to changes in likelihood* (Wakker, 2004).

In our model of non-rational Bayesian learning we restrict attention to a class of capacities that are defined as *neo-additive capacities* in the sense of Chateauneuf, Eichberger, and Grant (2006). Neo-additive capacities stand for deviations from additive probabilities such that a parameter δ (*degree of ambiguity*) measures the lack of confidence the decision maker has in some subjective additive probability distribution π. In addition, a second parameter λ measures the degree of optimism versus pessimism by which the decision maker resolves his ambiguity.

Definition: *A* neo-additive capacity, *ν, is defined, for some $\delta, \lambda \in [0,1]$, by*

$$\nu(A) = \delta \cdot (\lambda \cdot \omega^o(A) + (1-\lambda) \cdot \omega^p(A)) + (1-\delta) \cdot \pi(A) \tag{1}$$

[3] CEU theory was first axiomatized by Schmeidler (1986, 1989) within the Anscombe and Aumann (1963) framework, which assumes preferences over objective probability distributions. Subsequently, Gilboa (1987) as well as Sarin and Wakker (1992) have presented CEU axiomizations within the Savage (1954) framework, assuming a purely subjective notion of likelihood. CEU theory is equivalent to *cumulative prospect theory* (Tversky and Kahneman, 1992; Wakker and Tversky, 1993) restricted to the domain of gains (compare Tversky and Wakker, 1995). Moreover, as a representation of preferences over lotteries, CEU theory coincides with *rank dependent utility theory* as introduced by Quiggin (1981, 1982).

for all $A \in \mathcal{B}$ such that π 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

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

Let the state space Ω be finite and denote by $f(\omega)$ the outcome of the Savage-act f in state $\omega \in \Omega$. The Choquet expected utility of a Savage-act f with respect to a neo-additive capacity is

$$CEU(f, \nu) = \delta \cdot \left(\lambda \cdot \max_{\omega \in \Omega} u(f(\omega)) + (1 - \lambda) \cdot \min_{\omega \in \Omega} u(f(\omega)) \right) + (1 - \delta) \cdot \sum_{\omega \in \Omega} \pi(\omega) \cdot u(f(\omega)), \tag{2}$$

with $u(\cdot)$ denoting von Neumann-Morgenstern utility indices.[4]

In contrast to EU preferences there exist for CEU preferences several possibilities of deriving from ex ante preferences ex post preferences, i.e., preferences conditional on the fact that some event B has occurred. Following Gilboa and Schmeidler (1993) we focus on so-called *f-Bayesian update rules* for preferences $\succeq$ over Savage acts. That is, we consider some collection of conditional preference orderings, $\left\{ \succeq_B^f \right\}$ for all events B, such that for all acts g, h

$$g \succeq_B^f h \Leftrightarrow (g, B; f, \neg B) \succeq (h, B; f, \neg B) \tag{3}$$

where $(g, B; f, \neg B)$ denotes the act that gives consequences $g(\omega)$ for all $\omega \in B$ and consequences $f(\omega)$ for all $\omega \in \neg B$. Gilboa and Schmeidler show that CEU preferences

[4] Ludwig and Zimper (2006a) show that the CEU of an act with respect to a neo-additive capacity can be equivalently described by the *α-maxmin expected utility with respect to multiple priors* (α-MEU) of an act which encompasses the original multiple priors approach of Gilboa and Schmeidler (1989) as a special case (see, e.g., Ghirardato et al., 1998; Ghirardato et al., 2004; Siniscalchi, 2005). In particular, we have equivalence between the CEU with respect to neo-additive capacities and the α-MEU with respect to so-called *ε-contaminated priors* used in Bayesian statistics (Berger and Berliner, 1986) that may be interpreted as neo-additive capacities.

$\succeq$ on Savage acts are updated to conditional CEU preferences $\left\{\succeq_B^f\right\}$ for all events B if and only if f is an act such that for some event $E \in \Omega$

$$f = (x^*, E; x_*, \neg E), \tag{4}$$

where x^* denotes the best and x_* denotes the worst consequence possible. The different possible specifications of E in (4) can result in a multitude of different f-Bayesian update rules if *dynamic consistency* is violated. For example, for the so-called *optimistic* update rule f is the constant act where $E = \emptyset$. That is, under the optimistic update rule the null-event 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 A, the decision maker implicitly assumes that had A not occurred, the worst possible outcome [...] would have resulted. In other words, the behavior given A [...] exhibits 'happiness' that A 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 obtains

$$\nu^{opt}(A \mid B) = \frac{\nu(A \cap B)}{\nu(B)}. \tag{5}$$

For the *pessimistic* (=Dempster-Shafer) update rule f is the constant act where $E = S$, associating with the null-event the best consequence possible. Gilboa and Schmeidler:

> "[...] 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

$$\nu^{pess}(A \mid B) = \frac{\nu(A \cup \neg B) - \nu(\neg B)}{1 - \nu(\neg B)}. \tag{6}$$

Observation 1: *Suppose* $A, B \notin \{\varnothing, \Omega\}$.

(i) *An application of the optimistic update rule (5) to a prior belief (1) results in the conditional belief*

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

with

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

(ii) *An application of the pessimistic update rule (6) to a prior belief (1) results in the conditional belief*

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

with

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

Proof: Relegated to the appendix.

Let $A, B \notin \{\varnothing, \Omega\}$ and observe that

$$\nu^{pess}(A \mid B) < \nu^{opt}(A \mid B), \tag{7}$$

if $\delta > 0$, and

$$\nu^{pess}(A \mid B) < \pi(A \mid B) < \nu^{opt}(A \mid B)$$

if $\delta > 0$, $\lambda \in (0,1)$, and $\pi(A \mid B) \in (0,1)$. For the ex post evaluation of any Savage act f we therefore have

$$CEU(f, \nu^{pess}(A \mid B)) \leq EU(f, \pi(A \mid B)) \leq CEU\left(f, \nu^{opt}(A \mid B)\right),$$

whereby these inequalities are strict in non-trivial cases.

3 Psychologically biased Bayesian learning

Let us first recall the standard model of *rational* Bayesian learning which obtains as a special case of our approach. Following Viscusi and O'Connor (1984) we consider an agent who has a prior probability distribution over the π^* parameter of a Binomial-distribution (π^* being the *true* probability of some event A) such that this prior distribution belongs to the family of Beta distributions. The agent's subjective prior about π^*, denoted π, is the expected value of this prior Beta distribution, i.e., $\pi(A) = \frac{\alpha}{\alpha+\beta}$ for given parameters α and β. Let $\pi(A \mid \mathcal{I}_n)$ denote the agent's posterior about π derived from rational Bayesian learning whereby the event $\mathcal{I}_n$ denotes information equivalent to a statistical experiment in which event A has occurred k_n-times in n independent trials.

Rational Bayesian learning then results in a posterior Beta distribution about π^* with expected value $\frac{\alpha+k_n}{\alpha+\beta+n}$ implying for the posterior belief

$$\pi\left(A \mid \mathcal{I}_n\right)=\left(\frac{\alpha+\beta}{\alpha+\beta+n}\right)\cdot\pi\left(A\right)+\left(\frac{n}{\alpha+\beta+n}\right)\cdot\mu_n \tag{8}$$

where μ_n denotes the sample mean $\frac{k_n}{n}$. Since $\frac{k_n}{n}$ converges for $n\rightarrow\infty$ in probability to $\pi^*\left(A\right)$, we have, for any $c>0$,

$$\lim_{n\rightarrow\infty} prob\left(\left|\pi\left(A\mid\mathcal{I}_n\right)-\pi^*\left(A\right)\right|<c\right)=1, \tag{9}$$

which we abbreviate henceforth as

$$\lim_{n\rightarrow\infty}\pi\left(A\mid\mathcal{I}_n\right)=\pi^*\left(A\right).$$

That is, under the assumption of rational Bayesian learning the posterior beliefs converges to the true probability if the number of trials (=sample size) approaches infinity.[5]

Definition: Psychologically biased Bayesian learning.

(i) *We say an agent is* optimistically biased *if his posterior beliefs result from an application of the optimistic Bayesian update rule (5) to a neo-additive prior (1) such that the additive part of the posterior reflects rational Bayesian learning in the sense of (8).*

(ii) *We say an agent is* pessimistically biased *if his posterior beliefs result from an application of the pessimistic Bayesian update rule (6) to a neo-additive prior (1) such that the additive part of the posterior reflects rational Bayesian learning in the sense of (8).*

Let us assume that an agent who receives information in the n-th trial, did also receive information in the proceeding trials; that is, the events $\mathcal{I}_n$, $n=1,2,..$, form a nested sequence $\mathcal{I}_1\supseteq\mathcal{I}_2\supseteq\ldots$. As a consequence the corresponding sequence of probabilities $\pi\left(\mathcal{I}_1\right),\pi\left(\mathcal{I}_2\right),\ldots$ is monotonically decreasing, implying the existence of a unique limit point, i.e.,

$$\lim_{n\rightarrow\infty}\pi\left(\mathcal{I}_n\right)=b\in\left[0,\pi\left(\mathcal{I}_1\right)\right]. \tag{10}$$

[5] A similar result obtains when the agent has a normally distributed prior over the mean of some normal distribution (cf. Tonks 1983).

Moreover, in the plausible case that the agent does not expect to collect new information forever we have $\lim_{n\rightarrow\infty}\pi\left(\mathcal{I}_n\right)=0$. Because of (9) and (10) we obtain for the limit beliefs of psychologically biased Bayesian learning:

Observation 2: *Suppose* $A\notin\{\varnothing,\Omega\}$.

(i) *If the agent is optimistically biased, then*

$$\lim_{n\rightarrow\infty}\nu^{opt}\left(A\mid\mathcal{I}_n\right)=\delta^{opt}+\left(1-\delta^{opt}\right)\cdot\pi^*\left(A\right)$$

such that

$$\delta^{opt}\in\left[\frac{\delta\cdot\lambda}{\delta\cdot\lambda+(1-\delta)\cdot\pi\left(\mathcal{I}_1\right)},1\right].$$

(ii) *If the agent is pessimistically biased, then*

$$\lim_{n\rightarrow\infty}\nu^{pess}\left(A\mid\mathcal{I}_n\right)=\left(1-\delta^{pess}\right)\cdot\pi^*\left(A\right)$$

such that

$$\left(1-\delta^{pess}\right)\in\left[0,\frac{(1-\delta)\cdot\pi\left(\mathcal{I}_1\right)}{\delta\cdot(1-\lambda)+(1-\delta)\cdot\pi\left(\mathcal{I}_1\right)}\right].$$

Observe that $\delta^{opt}>0$ if and only if $\delta>0$ and $\lambda>0$. Analogously, $\left(1-\delta^{pess}\right)<1$ if and only if $\delta>0$ and $\lambda<1$.

Corollary: *Suppose* $\pi^*\left(A\right)\in(0,1)$ *and* $\delta>0$.

(i) *The posteriors of an optimistically biased agent with* $\lambda>0$ *converge to some belief strictly greater than the true probability* $\pi^*\left(A\right)$. *In particular, if*

$$\lim_{n\rightarrow\infty}\pi\left(\mathcal{I}_n\right)=0,$$

then the agent's posteriors converge to extreme optimism, i.e.,

$$\lim_{n\rightarrow\infty}\nu^{opt}\left(A\mid\mathcal{I}_n\right)=1.$$

(ii) *The posteriors of a pessimistically biased agent with* $\lambda<1$ *converge to some belief strictly greater than the true probability* $\pi^*\left(A\right)$. *If*

$$\lim_{n\rightarrow\infty}\pi\left(\mathcal{I}_n\right)=0,$$

then the agent's posteriors converge to extreme pessimism, i.e.,

$$\lim_{n\rightarrow\infty}\nu^{pess}\left(A\mid\mathcal{I}_n\right)=0.$$

The corollary demonstrates that psychologically biased Bayesian learning in our sense violates the two standard paradigms of rational Bayesian learning. Firstly, the posterior "subjective" beliefs do not converge to the "objective" probabilities in an infinite learning process. Secondly, the posteriors of two different agents do not converge to the same limit belief if they receive the same information but interpret it differently.

4 A first result: Identical information partitions

Throughout this paper we adopt Aumann's (1976) original epistemic framework. We consider two partitions $\mathcal{P}_1$ and $\mathcal{P}_2$ of a non-empty state-space Ω which are interpreted as the information space of agent 1, respectively 2. Denote by $P_i(\omega)$, with $i \in \{1,2\}$, the member of $\mathcal{P}_i$ that contains $\omega \in \Omega$. We say that *i knows event $A \in \mathcal{B}$ in state ω* iff $P_i(\omega) \subseteq A$. Moreover, let $\mathcal{P}_1 \wedge \mathcal{P}_2$ denote the finest partition of $\mathcal{B}$ that is coarser than $\mathcal{P}_1$ and $\mathcal{P}_2$ (i.e., the *meet* of $\mathcal{P}_1$ and $\mathcal{P}_2$). Following Aumann's definition, we say that *event $A \in \mathcal{B}$ is common knowledge between agent 1 and 2 in state ω* iff $P(\omega) \subseteq A$ whereby $P(\omega)$ is the member of $\mathcal{P}_1 \wedge \mathcal{P}_2$ containing $\omega \in \Omega$.

Our first *agreeing to disagree* result considers the situation in which agents have identical information partitions but apply different update rules.

Proposition 1: *Consider the following assumptions:*

- **(A1)** *The agents have identical neo-additive priors, i.e., $\nu_1 = \nu_2 \equiv \nu$, such that $\delta > 0$.*
- **(A2)** *The agents have identical information partitions $\mathcal{P}_1 = \mathcal{P}_2 \neq \{\Omega\}$.*
- **(A3)** *Agent 1 is optimistically whereas agent 2 is pessimistically biased.*
- **(A4)** *The agents' posteriors are common knowledge in some state of the world $\omega^* \in \Omega$.*

Then the agents' posterior beliefs about any event $A \notin \{\varnothing, \Omega\}$ are different.

Proof: Suppose that the posteriors are common-knowledge in $\omega^* \in \Omega$. By assumption, agent 1 is optimistically and agent 2 is pessimistically biased, implying

$$\begin{aligned}\nu_1(A \mid P_1(\omega^*)) &= \nu^{opt}(A \mid P_1(\omega^*)) \\ \nu_2(A \mid P_2(\omega^*)) &= \nu^{pess}(A \mid P_2(\omega^*)).\end{aligned}$$

Moreover, $\mathcal{P}_1 = \mathcal{P}_2$ implies $\mathcal{P}_1 = \mathcal{P}_2 = \mathcal{P}_1 \wedge \mathcal{P}_2$ so that $P(\omega^*) = P_1(\omega^*) = P_2(\omega^*)$. By inequality (7), the agents' posteriors $\nu_1(A \mid P(\omega^*))$ and $\nu_2(A \mid P(\omega^*))$ are therefore different for every event $A \notin \{\varnothing, \Omega\}$.$\square$

Proposition 1 shows that, except for degenerate cases, optimistically and pessimistically biased agents have in the ex-post situation always strict incentives to bet with each other. By the corollary to observation 2, these incentives will get stronger the more information the agents collect. This result holds regardless of the fact whether the agents collect identical or different information.

While the formal proof of proposition 1 is simple, it reveals a fundamental difference between Aumann's concept of information and our approach. According to Aumann, any differences in the beliefs of different agents are caused by different information received by the agents. According to our approach, differences in the beliefs of agents may also result because of different psychological attitudes with respect to the *interpretation* of new information. That is, while one agent might have a "half-full" attitude, another agent may have a "half-empty" attitude when interpreting the same fact.

5 An illustrative example: Asset-trading

We illustrate proposition 1 by a simple example in which ex-post asset-trading happens in every state of the world due to different ex-post evaluations of the asset. Moreover, these different ex-post evaluations are common knowledge to the agents despite the fact that their ex ante evaluations and their information partitions are identical.

Assume that agent 2 owns in period 1 a financial asset which gives vNM utility of 1 in case an investment project is successful and an utility of 0 in case it is not. Before it will be revealed in period 3 whether the project is successful or not, there will be news about the project's progress, either *good* or *bad*, in period 2. Let the relevant state space be given as

$$\Omega = \{SG, SB, FG, FB\}$$

whereby the event $G = \{SG, FG\}$ stands for *good* and the event $B = \{SB, FB\}$ stands for *bad* news in period 2. Accordingly, $S = \{SG, SB\}$ is the event of *success* and $F = \{FG, FB\}$ is the event of *failure*. The information partitions $\mathcal{P}_1(t), \mathcal{P}_2(t)$, $t \in \{1, 2\}$, in period $t = 1$ are

$$\mathcal{P}_1(1) = \mathcal{P}_2(1) = \{\Omega\}.$$

Under the assumption of identical neo-additive priors $\nu_1 = \nu_2 = \nu$, both agents therefore (ex-ante) evaluate the Savage-act f of *holding the asset* by the same CEU (2), namely,

$$CEU_1(f,\nu) = CEU_2(f,\nu) = \delta \cdot \lambda + (1-\delta) \cdot \pi(S).$$

As a consequence, there is no strict incentive for the agents to trade the asset in the ex-ante situation.

Consider now the following information partitions at period 2

$$\mathcal{P}_1(2) = \mathcal{P}_2(2) = \{\{SG, FG\}, \{SB, FB\}\}$$

and assume that agent 1 applies optimistically and agent 2 applies pessimistically biased Bayesian learning upon learning the news $x \in \{G, B\}$. Agent 1, resp. 2, then evaluate *holding the asset* in the ex-post situation as

$$CEU_1\left(f, \nu^{opt}(\cdot \mid x)\right) = \nu^{opt}(S \mid x),$$

resp.

$$CEU_2\left(f, \nu^{pess}(\cdot \mid x)\right) = \nu^{pess}(S \mid x).$$

By (7), we have for $\delta, \lambda \in (0,1)$

$$CEU_2\left(f, \nu^{pess}(\cdot \mid x)\right) < \pi(S \mid x) < CEU_1\left(f, \nu^{opt}(\cdot \mid x)\right).$$

Thus, regardless of whether the news turn out *good* or *bad* agent 1 ex-post evaluates the asset strictly higher than agent 2. As a consequence there will be ex-post trade in the asset in every state of the world. For example, at price $\pi(S \mid x)$ agent 2 would strictly prefer to sell the asset while agent 1 would strictly prefer to buy it.

Remark. Observe that if both agents were EU decision makers, i.e., $\delta = 0$, there would be no strict incentive for ex-post trading if there is no strict incentive for ex-ante trading. As this example shows, this is not necessarily true for CEU decision makers because of the possibility of dynamically inconsistent CEU preferences.[6] According to our concept of psychologically biased Bayesian learning, the incentive for ex-post trading results in the example from the agents' different psychological attitudes with respect to the interpretation of new information.

[6] Ludwig and Zimper (2006b) demonstrate that *sophisticated* (Strotz 1956, Pollak 1968) CEU decision makers may have even stronger incentives for *intrapersonal commitment* than *sophisticated hyperbolic discounting* decision makers in the sense of Laibson (1997) and Frederick, Loewenstein, and O'Donoghue (2002).

6 A second result: Identical learning rules

Our second *agreeing to disagree* result applies to people who use the same psychologically biased learning rule but have different information partitions.

Proposition 2: *Consider the following assumptions:*

(A1') *The agents have identical neo-additive priors, i.e.,* $\nu_1 = \nu_2 \equiv \nu$*, such that* $\delta > 1$.

(A2') *Both agents are either optimistically or pessimistically biased.*

(A3') *The agents' posteriors are common-knowledge in some state of the world* $\omega^* \in \Omega$.

(A4') *The agents' priors satisfy* $\nu(P_1(\omega^*)) \neq \nu(P_2(\omega^*))$ *whereby* $P_1(\omega^*), P_2(\omega^*) \neq \Omega$.

Then the agents' posterior beliefs about any event $A \notin \{\varnothing, \Omega\}$ *are different.*

Observe that assumption (A4'), i.e., $\nu(P_1(\omega^*)) \neq \nu(P_2(\omega^*))$, cannot hold if the agents have identical priors **and** identical information partitions. That is, the result of proposition 2 only applies in situations of asymmetric information, i.e., $\mathcal{P}_1 \neq \mathcal{P}_2$, such that the two events $P_1(\omega^*)$ and $P_2(\omega^*)$ are not equally likely according to the agents' common prior.

Before we turn to the proof of proposition 2 consider the following example which illustrates the intuition behind our formal proof.

Example. Consider the following information structure

$$\begin{aligned} \mathcal{P}_1 &= \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}, ...\} = \{P_1^1, P_1^2, ...\}, \\ \mathcal{P}_2 &= \{\{\omega_1, \omega_2, \omega_3, \omega_4\}, ...\} = \{P_2^1, ...\} \end{aligned}$$

so that

$$\mathcal{P}_1 \wedge \mathcal{P}_2 = \{\{\omega_1, \omega_2, \omega_3, \omega_4\}, ...\}.$$

Suppose agent 1 and 2 have a common neo-additive prior ν with $\delta \in (0,1)$ such that

$$\pi(\{\omega_1\}) = ... = \pi(\{\omega_4\}) > 0.$$

Further suppose that both agents are optimistically biased. Let

$$A = \{\omega_2, \omega_3\}$$

and observe that

$$\nu_1\left(A \mid P_1^1\right) = \nu_1\left(A \mid P_1^2\right) = \delta_1^{opt} + \left(1-\delta_1^{opt}\right)\cdot\frac{1}{2} \tag{11}$$

with

$$\delta_1^{opt} = \frac{\delta\cdot\lambda}{\delta\cdot\lambda + (1-\delta)\cdot\pi\left(\{\omega_1,\omega_2\}\right)}$$

and

$$\nu_2\left(A \mid P_2^1\right) = \delta_2^{opt} + \left(1-\delta_2^{opt}\right)\cdot\frac{1}{2} \tag{12}$$

with

$$\delta_2^{opt} = \frac{\delta\cdot\lambda}{\delta\cdot\lambda + (1-\delta)\cdot\pi\left(\{\omega_1,\omega_2,\omega_3,\omega_4\}\right)}.$$

Observe that the posterior of each agent is the same in every state belonging to $P(\omega^*) \in \mathcal{P}_1 \wedge \mathcal{P}_2$ with $\omega^* \in \{\omega_1,\omega_2,\omega_3,\omega_4\}$ so that we can stipulate that the agents' posteriors are common knowledge in every state $\omega^* \in \{\omega_1,\omega_2,\omega_3,\omega_4\}$. Since $\delta > 0$, the posteriors (11) and (12) coincide if and only if

$$\begin{aligned}
\delta_1^{opt} &= \delta_2^{opt} \Leftrightarrow \\
\pi\left(\{\omega_1,\omega_2\}\right) &= \pi\left(\{\omega_1,\omega_2,\omega_3,\omega_4\}\right) \Leftrightarrow \\
\nu\left(P_1\left(\omega^*\right)\right) &= \nu\left(P_2\left(\omega^*\right)\right),
\end{aligned}$$

which is not the case in this example. Thus, despite identical priors and identical Bayesian learning rules, both agents have different posterior beliefs which are common knowledge.

Proof of proposition 2. Our proof builds on Aumann's (1976) original proof for the case of an additive probability measure, i.e., $\delta = 0$.

Step 1. Aumann (1976): For an additive common prior π the agents' posteriors must be identical when they are common knowledge at some state of the world.

Suppose to the contrary that there is some $\omega^* \in \Omega$ in which it is common knowledge that

$$\pi_1\left(A \mid P_1\left(\omega^*\right)\right) = q_1 \text{ and } \pi_2\left(A \mid P_2\left(\omega^*\right)\right) = q_2$$

such that $q_1 \neq q_2$ for some event $A \in \mathcal{B}$. Then

$$\pi_1\left(A \mid P_1^j\right) = q_1, \tag{13}$$

for all $P_1^j \subseteq P(\omega^*)$ whereby $P(\omega^*)$ is the member of $\mathcal{P}_1 \wedge \mathcal{P}_2$ containing ω^*. Denote by $P_1^1, ..., P_1^n$ the members of $\mathcal{P}_1$ such that

$$P_1^1 \cup ... \cup P_1^n = P\left(\omega^*\right) .$$

By additivity,

$$\pi\left(P_1^1\right)+...+\pi\left(P_1^n\right)=\pi\left(P\left(\omega^*\right)\right) \tag{14}$$

since $P_1^1,...,P_1^n$ is a partition of $P(\omega^*)$. Also by additivity,

$$\pi_1\left(A\mid P_1^j\right)=\frac{\pi\left(A\cap P_1^j\right)}{\pi\left(P_1^j\right)},\ j=1,..,n$$

so that, by (13),

$$\pi\left(P_1^1\right)+...+\pi\left(P_1^n\right)=\frac{\pi\left(A\cap P_1^1\right)}{q_1}+...+\frac{\pi\left(A\cap P_1^n\right)}{q_1}.$$

Since, by additivity,

$$\pi\left(A\cap P_1^1\right)+...+\pi\left(A\cap P_1^n\right)=\pi\left(A\cap P\left(\omega^*\right)\right)$$

we have

$$\pi\left(P_1^1\right)+...\pi\left(P_1^n\right)=\frac{\pi\left(A\cap P\left(\omega^*\right)\right)}{q_1}.$$

Thus, by (14),

$$\frac{\pi\left(A\cap P\left(\omega^*\right)\right)}{q_1}=\pi\left(P\left(\omega^*\right)\right).$$

An analogous argument for agent 2 results in

$$\frac{\pi\left(A\cap P\left(\omega^*\right)\right)}{q_2}=\pi\left(P\left(\omega^*\right)\right)$$

implying the desired result $q_1=q_2$.$\square$

Step 2. Consider now the case of identical non-additive priors (1), i.e., $\delta>0$. Let $A\notin\{\varnothing,\Omega\}$ and suppose both agents are optimistically biased; (there is an analogous argument for pessimistically biased agents). Then, for $\omega\in\Omega$,

$$\nu_1^{opt}\left(A\mid P_1\left(\omega\right)\right)=\delta_{1,B}^{opt}+\left(1-\delta_{1,B}^{opt}\right)\cdot\pi\left(A\mid P_1\left(\omega\right)\right) \tag{15}$$

with

$$\delta_{1,B}^{opt}=\frac{\delta\cdot\lambda}{\delta\cdot\lambda+(1-\delta)\cdot\pi\left(P_1\left(\omega\right)\right)}$$

and

$$\nu_2^{opt}\left(A\mid P_2\left(\omega\right)\right)=\delta_{2,B}^{opt}+\left(1-\delta_{2,B}^{opt}\right)\cdot\pi\left(A\mid P_2\left(\omega\right)\right) \tag{16}$$

with

$$\delta_{2,B}^{opt}=\frac{\delta\cdot\lambda}{\delta\cdot\lambda+(1-\delta)\cdot\pi\left(P_2\left(\omega\right)\right)}.$$

Assume now that the posteriors (15) and (16) are common knowledge in some state $\omega^* \in \Omega$. By the argument of step 1, the posteriors must coincide for the special case of $\delta = 0$. We therefore have for the additive part of the posteriors

$$\pi\left(A \mid P_1\left(\omega^*\right)\right) = \pi\left(A \mid P_2\left(\omega^*\right)\right),$$

so that the agents' posteriors (15) and (16) are different if and only if

$$\begin{aligned} \delta_{1,B}^{opt} &\neq \delta_{2,B}^{opt} \Leftrightarrow \\ \pi\left(P_1\left(\omega^*\right)\right) &\neq \pi\left(P_2\left(\omega^*\right)\right) \Leftrightarrow \\ \nu\left(P_1\left(\omega^*\right)\right) &\neq \nu\left(P_2\left(\omega^*\right)\right). \end{aligned}$$

This proves the proposition. □□

7 Concluding remarks

In a first step, we have developed a model of psychologically biased Bayesian learning whereby we focus on the two benchmark cases of optimistically, resp. pessimistically, biased learning. While our model encompasses the standard model of rational Bayesian learning as a special case, it additionally allows for the possibility that an agent exhibits a "myside bias" in the interpretation of new information.

In a second step, we apply our model of psychologically biased Bayesian learning to the epistemic situation studied in Aumann (1976). Two main results emerge:

1. Even if people receive the same information, they may agree to disagree if their psychologically attitudes about the interpretation of new information are different.

2. Even if people have the same psychologically attitudes, they may agree to disagree if they receive different information.

Both results are in contrast to Aumann's famous conclusion that agents cannot agree to disagree regardless of whether they receive the same information or not. Our concept of psychologically biased Bayesian learning can therefore offer a possible explanation for the existence of ex-post trade in financial assets.

Appendix

Proof of observation 1:

Applying the optimistic Bayesian update rule to a neo-additive capacity gives, for $A \notin \{\varnothing, \Omega\}$,

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

such that

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

Applying the pessimistic Bayesian update rule to a neo-additive capacity gives, for $A \notin \{\varnothing, \Omega\}$,

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

such that

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

□

References

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

Aumann, R. (1976), "Agreeing to disagree," *Annals of Statistics* **4**, 1236-1239.

Bacharach, M. (1985), "Some Extension of a Claim of Aumann in an Axiomatic Model of Knowledge", *Journal of Economic Theory* **37**, 167–190.

Baron, J. (2000), *Thinking and Deciding*, Cambridge University Press: New York, Melbourne, Madrid.

Beck, A.T. (1976), *Cognitive Therapy and the Emotional Disorders*, International Universities Press: New York.

Berger, J., and L.M. Berliner (1986), "Robust Bayes and Empirical Bayes Analysis with ε-Contaminated Priors", *Annals of Statistics* **14**, 461-486.

Bonanno, G., and K. Nehring (1999), "How to Make Sense of the Common Prior Assumption under Incomplete Information", *International Journal of Game Theory* **28**, 409-434.

Chateauneuf, A., Eichberger, J., and S. Grant (2006), "Choice under Uncertainty with the Best and Worst in Mind: Neo-additive Capacities", *Journal of Economic Theory* forthcoming.

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

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 Expected Utility Preferences", mimeo

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 M. Le Breton (1993), "Dynamically Consistent Beliefs Must Be Bayesian", *Journal of Economic Theory* **61**, 1-22.

Frederick, S., Loewenstein, G., and T. O'Donoghue (2002), "Time Discounting and Time Preference: A Critical Review", *Journal of Economic Literature* **Vol. XL**, 351-401.

Geanakoplos, J. (1992), "Common Knowledge", *Journal of Economic Perspectives* **6**, 53–82.

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

Ghirardato, P., Klibanoff, P., and M. Marinacci (1998), "Additivity with Multiple Priors", *Journal of Mathematical Economics* **30**, 405-420.

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 (1993), "Updating Ambiguous Beliefs", *Journal of Economic Theory* **59**, 33-49.

Halevy, Y. (1998), "Trade between Rational Agents as a Result of Asymmetric Information", mimeo.

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

Harsanyi, J.C. (1967), "Games with Incomplete Information Played by 'Bayesian' Players. Part I: The Basic Model", *Management Science* **14**, 159-182.

Kahneman, D., and A. Tversky (1979), "Prospect theory: An Analysis of Decision under Risk", *Econometrica* **47**, 263-291.

Laibson, D. (1997), "Golden Eggs and Hyperbolic Discounting", *Quarterly Journal of Economics* **112**, 443-477.

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 (2006a), "Rational Expectations and Ambiguity: A Comment on Abel (2002)", *Economics Bulletin*, **4**(2), 1-15.

Ludwig, A., and A. Zimper (2006b), "Investment Behavior under Ambiguity: The Case of Pessimistic Decision Makers," *Mathematical Social Sciences* **52**, 111-130.

Milgrom, P. (1981), "An Axiomatic Characterization of Common Knowledge", *Econometrica* **49**, 219-222.

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.

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.

Pollak, R.A. (1968), "Consistent Planning", *The Review of Economic Studies* **35**, 201-208.

Quiggin, J.P. (1981), "Risk Perception and Risk Aversion among Australian Farmers", *Australian Journal of Agricultural Economics* **25**, 160-169.

Quiggin, J.P. (1982), "A Theory of Anticipated Utility", *Journal of Economic Behavior and Organization* **3**, 323-343.

Rubinstein, A., and A. Wollinski (1990), "On the Logic of "Agreeing to Disagree" Type Results", *Journal of Economic Theory* **51**, 184-193.

Samet, D. (1990), "Ignoring Ignorance and Agreeing to Disagree", *Journal of Economic Theory* **52**, 190-207.

Sarin, R., and P.P. Wakker (1992), "A Simple Axiomatization of Nonadditive Expected Utility", *Econometrica* **60**, 1255-1272.

Sarin, R., and P.P. Wakker (1998a), "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. (2001), "Bayesian Updating for General Maxmin Expected Utility Preferences", mimeo.

Siniscalchi, M. (2005), "A Behavioral Characterization of Plausible Priors", *Journal of Economic Theory*, forthcoming.

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

Strotz, R.H. (1956), "Myopia and Inconsistency in Dynamic Utility Maximization", *The Review of Economic Studies* **23**, 165-180.

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

Tversky, A., and D. Kahneman (1992), "Advances in Prospect Theory: Cumulative Representations of Uncertainty", *Journal of Risk and Uncertainty* **5**, 297-323.

Tversky, A., and P.P. Wakker (1995), "Risk Attitudes and Decision Weights", *Econometrica* **63**, 1255-1280.

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, and A. Tversky (1993), "An Axiomatization of Cumulative Prospect Theory", *Journal of Risk and Uncertainty* **7**, 147-176.

Yaari, M.E. (1987), "The Dual Theory of Choice under Risk", *Econometrica* **55**, 95-115.