



**DEPARTMENT OF INTERNATIONAL AND
EUROPEAN ECONOMIC STUDIES**

ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

**AMBIGUITY AVERSION, MODERN
BAYESIANISM AND SMALL WORLDS**

PHOEBE KOUNDOURI

NIKITAS PITTIS

PANAGIOTIS SAMARTZIS

NIKOLAOS ENGLEZOS

ANDREAS PAPANDREOU

Working Paper Series

17-07

November 2017

Ambiguity Aversion, Modern Bayesianism and Small Worlds

Phoebe Koundouri* Nikitas Pittis[†] Panagiotis Samartzis[‡]
Nikolaos Englezos[§] Andreas Papandreou[¶]

Abstract

1 Introduction

Since 1961, year which Ellsberg introduced the concept of "ambiguity aversion" (AA) in the form of his well-known thought experiments (see Ellsberg 1961), economists, psychologists and decision theorists have been trying to provide answers to the following two questions: (i) is AA empirically documented by means of formal experimentation and (ii) does AA violate the standard Bayesian conditions of rationality, or is it consistent with rationality when properly modified?

With respect to the first question, Machina and Siniscalchi (2014) offer an extensive survey of the empirical literature (including studies from the insurance and medical literature) that spans over a period of 50 years. The results of these experiments lend support to AA, although some results that point towards "ambiguity neutral" or even "ambiguity seeking" behavior do exist (see, for example, Viscusi and Chesson 1999).

*School of Economics, Athens University of Economics and Business, Greece; Grantham Research Institute, London School of Economics and Political Science, UK; International Center for Research on the Environment and the Economy, Greece; e-mail: pkoundouri@aueb.gr (corresponding author)

[†]Department of Banking and Financial Management, University of Piraeus, Greece; International Center for Research on the Environment and the Economy, Greece; e-mail: npittis@unipi.gr

[‡]Department of Banking and Financial Management, University of Piraeus, Greece; School of Economics, Athens University of Economics and Business, Greece; International Center for Research on the Environment and the Economy, Greece; e-mail: psamartzis@unipi.gr

[§]Department of Banking and Financial Management, University of Piraeus, Greece; International Center for Research on the Environment and the Economy, Greece; e-mail: englezos@unipi.gr

[¶]School of Economics, National and Kapodistrian University of Athens, Greece; International Center for Research on the Environment and the Economy, Greece; e-mail: aap@econ.uoa.gr

As far as the second question is concerned, the main argument in favor of AA being consistent with rational probabilistic beliefs is the following: AA implies that probabilistic beliefs are not sophisticated, in the sense that they cannot be represented by a unique *prior* probability measure. Is a probabilistically non-sophisticated agent necessarily irrational? The advocates of the AA say no. Instead, they argue that the agent might possess so little (if any) information about the events she is interested in, that she is justifiably (and thus rationally) unable to form a unique prior distribution. As Gilboa and Schmeidler (1989) put it "...the subject has too little information to form a prior. Hence (s)he considers a set of priors as possible." (1989, pp. 142).

Schmeidler (1989) is one of the earliest and best-known attempts to provide an axiomatization of AA. Schmeidler argues that at the heart of AA lies the following asymmetry: Consider an event space \mathcal{F} which contains the "events of interest" for agent X. Assume that $\mathcal{F}_k \subset \mathcal{F}$ contains those events whose probabilities are known to X, whereas $\mathcal{F}'_k \subset \mathcal{F}$, ($\mathcal{F}_k \cap \mathcal{F}'_k = \emptyset$, $\mathcal{F}_k \cup \mathcal{F}'_k = \mathcal{F}$) contains the rest of events for which X has no information about. AA means that X favors betting on events in \mathcal{F}_k than \mathcal{F}'_k . However, under the maintained assumption that X's subjective probability of any event $D \in \mathcal{F}$ is measured by X's willingness to bet on D , AA implies that X's degree of belief defined on \mathcal{F} violate the additivity axiom of probability theory, thus X is not probabilistically sophisticated.

What is the main reason for the violation of the additivity property in X's degrees of belief, denoted by P_X ? The answer is that X has specific information I_S , which is relevant for some specific events in \mathcal{F} and irrelevant for all the rest. Is I_S allowed to be used in the determination of the prior beliefs, or should these beliefs be based exclusively on the so-called "background" information, I_B ?

How is I_B defined? A universally accepted sharp definition of I_B is hard to find. I_B is usually meant to include the "corpus of background knowledge" (Easwaran 2008) that describes the general characteristics of the chance set up in hand (e.g. there are n coins involved, each one is two-sided, successive tosses are independent, etc). More importantly, I_B should include the full set of hypotheses H that describe the probabilistic properties of the phenomenon of interest, the full set of evidential sentences E that are relevant for H together with the entailment relationships between elements of H and elements of E . On the other hand, I_B should not include any actual outcomes of the experiment or information about the objective chances (e.g. coin is fair). This information is specific information or relevant evidence and should be thought of as part of I_S .

The question of whether to allow both I_B and I_S or only I_B to affect the generation of the agent's prior belief function is crucial for deciding whether AA is (or can be made) consistent with rationality. Hence the crucial question is: should P_0 depend only on I_B (Option I) or on both I_B and I_S (Option II)? There are two approaches in the literature: one that claims that the distinction between I_B and I_S is not only useful but compelling, and the other which sees no point in distinguishing between the two. As Howson and Urbach (1991) put it, Option II means that "the priors in any calculation can also be posterior

probabilities calculated from the historical data" (1991, pp. 374). As such, current probabilities (credences), P_t , are shaped by all the available information that the agent possesses at t , including both I_B and I_S . In such a case, the evidential significance of I_S is "built into" P_t . Notationally, we may emphasize this, by replacing P_t with $P_t^{I_B, I_S}$. Following Meacham (2007), we refer to this view as "Classical Bayesianism" (CB). On the other hand, Option I, hereafter referred to as "Modern Bayesianism" (MB), amounts to the prior, P_0 , being determined at a time in which no specific evidence I_S is available, meaning that X is capable of counterfactual probabilistic reasoning. Put differently, the prior probability function should be determined, once and for all, at a unique point in time, say t_0 , in which only I_B is available. At t_0 , I_S is contingent (hypothetical) rather than actual. Notationally, P_0 ¹ may be replaced by $P_0^{I_B}$.

The central question of this paper is whether AA is consistent with rationality and whether the distinction between CB and MB can provide insights to this question. The answer to this question is emphatically affirmative. The main point of this paper is to show that an MB agent, operating in the context of a "small world" (more on this below), cannot exhibit AA.

We also ask the question of whether MB is a normatively appealing hypothesis. To this end, we put forward three arguments in favor of MB, analyzed in detail in the section 3. In summary, we argue that an MB agent (as opposed to a CB one) (i) protects herself from falling into logical contradictions akin to the relation between her subjective probability and objective chance, (ii) fares better in cases of "old evidence" and (iii) allows psychological detachment from actual evidence.

What about MB's descriptive status? Is MB reasonably realistic as a model of actual behavior? Or is it the case that MB puts exceedingly high standards on the agent's probabilistic reasoning abilities? To answer this question, we must first define the context within which the agent operates. To this end, we distinguish between "small" and "large" worlds, a distinction already made by Savage (1954). For the purposes of the present paper, a "world" is defined to be the domain (language) \mathcal{L} of the agent's probabilistic beliefs. This domain includes all the hypotheses of interest H_i , all the relevant evidential statements E_j , together with the corresponding (logical) entailment relations between H_i and E_j . A world is small if (a) the number of hypotheses in \mathcal{L} is small and (b) \mathcal{L} does not evolve over time. These two assumptions imply that the agent is able to conceive in advance all possible hypotheses concerning the matter of interest. In this paper we argue that MB is a reasonable hypothesis in small worlds.

On the contrary, MB is not likely to work in cases in which a) \mathcal{L} is exceedingly rich and b) \mathcal{L} changes over time due to the formation of new concepts (hypotheses) by the agent. With respect to the second case, imagine a situation in which the agent's current domain, \mathcal{L}_t , is richer than her domain \mathcal{L}_0 at the beginning of her investigations, t_0 (that is before any specific information I_S is available). This may happen if between the time periods t_0 and t , a new hypothesis H_{new} was conceived by the agent. In such a case, H_{new} is part of

¹We let $P_0 \equiv P_{t_0}$, to avoid the latter, more burdensome, notation.

\mathcal{L}_t but not of \mathcal{L}_0 . MB requires the agent to "go back in time" in order to find herself in the epistemic state that corresponds to t_0 . Unfortunately, at this state, the agent is unable to find any probabilistic assignments to H_{new} simply because H_{new} was not included in \mathcal{L}_0 . In such a case, MB requires the agent to contemplate her probabilistic assignments not in the "old" domain \mathcal{L}_0 (which did not include H_{new}) but in the "new", expanded domain \mathcal{L}_t . However, this shift between domains means that the agent is forced to change all her probabilistic assignments made in the context of \mathcal{L}_0 . This is the point at which MB loses its normative and descriptive edge. As Bammer and Smithson (2012, pp 95) remark "...changing the language (domain) in which items of evidence and hypotheses are expressed will typically change the confirmation relations between them."

Do MB's losses translate to CB's gains? The answer is negative. Both MB and CB are equally vulnerable to the presence of large and evolving worlds because under those worlds, both MB and CB require frequent changes in prior probabilities P_0^{IB} and $P_t^{IB,IS}$. Such changes invalidate many of the most interesting Bayesian theorems, including the celebrated "washing out of priors"² result.

It is worth emphasizing that standard Bayesianism, in either MB or CB form, does not allow the formation of new hypotheses after the prior probability function has been formed. Instead, it assumes that the agent is able to think in advance of all the possibilities in the hypothesis space; in short, the agent is assumed to be "logically omniscient". Logical omniscience is considered by many critics as the Achilles heel of Bayesianism: it implies an all-knowing agent who is able to track down all the logical implications in \mathcal{L} . Obviously, the extent to which such an assumption is unrealistic depends on the degree of richness of the underlying domain \mathcal{L} . For simple domains, it is reasonable to assume that the agent can comprehend the few logical entailments that these domains contain, which renders standard Bayesianism realistic in small worlds.

The main points of the paper are the following: (i) An MB agent operating in small worlds does not exhibit AA. (ii) AA may be thought of as a (potential) property of large worlds, because in such worlds MB is likely to be infeasible. (iii) The motivation of AA in the form of examples such as Ellsberg's classic urn-based or Schmeidler's coin-based ones is clearly poor, since these examples are textbook cases of small worlds (or elementary domains), in which the normatively preferable MB strategy is attainable.

The paper is organized as follows: Section 2 describes Schmeidler's two-coin thought experiment that is supposed to produce AA. It also defines formally two of the central concepts of this paper, namely MB and CB. Section 3 is an analytic discussion of the arguments for preferring MB over CB. Section 4 provides the main objection to MB, namely logical omniscience. It also elaborates on the connection between MB and small worlds and the efficiency of counterfactual probabilistic reasoning in such worlds. Section 5, demonstrates formally

²Difference in priors doesn't matter since in the long run, as more evidence is accumulated, these priors "wash out" in the sense that posterior subjective probabilities merge to the objective one.

the absence of AA under MB reasoning within Schmeidler's two-coin example. Section 6 concludes the paper.

2 Schmeidler's Two-Coin Example and Modern versus Classical Bayesianism

Schmeidler (1989) uses the following coin example, which aims at conveying the same message with Ellsberg's two-urn paradox (1961). Assume that the agent X considers two coins A and B. X knows with certainty that the probability of "heads" for A (H_A) is 0.5. On the other hand, she has no such information about B. The two coins are about to be tossed and X has the option to bet either on A-related events or B-related events. Specifically, she may choose bet A in which she earns 1\$ if the event

$$D_H^A = \{\text{coin A comes up heads}\} = \{H_A H_B, H_A T_B\}$$

occurs and loses 1\$ if the event

$$D_T^A = \{\text{coin A comes up tails}\} = \{T_A H_B, T_A T_B\}.$$

Alternatively, she is offered the option to enter bet B which gives her 1\$ if the event

$$D_H^B = \{\text{coin B comes up heads}\} = \{H_A H_B, T_A H_B\}$$

materializes and produces a loss of 1\$ if the event

$$D_T^B = \{\text{coin B comes up tails}\} = \{H_A T_B, T_A T_B\}$$

takes place. Obviously,

$$D_H^A, D_T^A \in \mathcal{F}_k$$

whereas

$$D_H^B, D_T^B \in \mathcal{F}'_k.$$

Indeed, the agent knows that

$$P(D_H^A) = P(D_T^A) = 0.5$$

although she is not informed about the probabilities $P(D_H^B)$ and $P(D_T^B)$. However, although she does not know the exact values of $P(D_H^B)$ and $P(D_T^B)$ she feels more willing to bet on D_H^A than D_H^B and on D_T^A than D_T^B , i.e.

$$P(D_H^B) < P(D_H^A)$$

and

$$P(D_T^B) < P(D_T^A).$$

It is easy to show that such a P is non-additive. Indeed,

$$D_H^B \cap D_T^B = \emptyset$$

and

$$D_H^B \cup D_T^B = \Omega$$

where Ω is the relevant sample space, namely

$$\Omega = \{H_A H_B, H_A T_B, T_A H_B, T_A T_B\}.$$

Assuming that P is additive,

$$1 = P(D_H^B \cup D_T^B) = P(D_H^B) + P(D_T^B) < 1$$

which is a contradiction.

What causes the violation of the additivity property in P ? It is not merely the presence of asymmetric information, but rather the fact that the agent allowed I_S to affect directly her probabilistic beliefs, instead of utilizing (as she should) I_S indirectly by conditionalization. But in order to be able to conditionalize on I_S , a pre-existing vehicle is required, namely a probability function, prior to the acquisition of I_S . Put differently, the agent should form her probabilistic beliefs by employing the MB counterfactual strategy briefly outlined in the introduction.

Let us now define the concepts of MB and CB more formally. To this end, assume an agent, who started her investigations at time t , having both I_B and I_S at her disposal. The agent is interested in updating her beliefs about the proposition/event A in the light of new evidence E obtained at $t + 1$. A CB agent carries out this task as follows:

Classical Bayesianism (CB):

$$P_{t+1, I_B, I_S, E}(A) = P_t^{I_B, I_S}(A | E). \quad (1)$$

On the other hand, the MB strategy advises the agent to pursue the following two-step procedure: First, "travel back" to your epistemic history and identify point, t_0 , at which I_S was not yet available (although it was conceivable). Second, once this point was found, evaluate the prior that you would have formed at this point (counterfactual or hypothetical prior) and use it from this point onwards as your actual prior.

Modern Bayesianism (MB):

$$P_{t+1, I_B, I_S, E}(A) = P_0^{I_B}(A | I_S, E). \quad (2)$$

Easwaran (2008) refers to $P_0^{I_B}$ "as expressing counterfactual possibilities of what she (the agent) imagines she would believe if she didn't have some of the information (in our case, I_S) she does now." (2008, pp.147). In a similar fashion, Meacham (2008) explains the epistemic role that $P_0^{I_B}$ is supposed to play: "A rational subject's credences are fixed by her hypothetical priors and her total evidence. A subject's credences are represented by a dynamic probability function, a function that changes with her evidence. A subject's hypothetical priors are represented by a static probability function, a function that encodes her disposition to respond to evidence. (Hypothetical priors are called 'priors' because

they can be thought of as a rational subject's original credences in possibilities, prior to the receipt of any evidence, and 'hypothetical' because it is unlikely that one ever was in such a state.)" (2008, emphasis added). It is worth emphasizing that despite the term "modern", MB is not a form of Bayesianism that was put forward only recently. Indeed, several authors from the 1970s entertained the idea of "hypothetical priors", that is, "priors without evidence"³.

What should the basic probability concept be in the development of rational decision theory? Prior or Current Probability? To answer this question we must think of what personal probabilities should reflect if they were to be rational: the person's permanent dispositions for forming beliefs (on the basis of the available evidence) or merely her momentary inclinations at some given point in time?

Rudolph Carnap argues strongly in favor of the first option. Indeed, his monumental work on inductive logic (Carnap 1950) is based on the concept of hypothetical or counterfactual initial credence function that can be ascribed to the agent X, before the collection of any evidence. More specifically, Carnap (1962) considers a sequence of data E_1, E_2, \dots, E_n obtained by the agent X up to the present time T_n (in Carnap's notation). He also defines K_n to be the "total observational knowledge" of X at T_n , that is

$$K_n = \cap_{i=1}^n E_i. \quad (3)$$

Then Carnap contemplates X's epistemic state at time T_0 in which X possesses no observational knowledge at all: "Now consider the sequence of X's credence functions. In the case of a human being we would hesitate to ascribe to him a credence function at a very early time point, before his abilities of reason and deliberate action are sufficiently developed. But again we disregard this difficulty by thinking either of an idealized human baby or of a robot. We ascribe to him a credence function Cr_1 for the time point T_1 ; Cr_1 represents X's personal probabilities based upon the datum E_1 as his only experience. *Going even one step further*, let us ascribe to him an initial credence function Cr_0 for the time point T_0 *before he obtains his first datum* E_1 . Any later function

³For example, Hesse (1975) suggests that the agent's prior probability should include only "previous background evidence which *does not enter the current calculations*" (Hesse, 1975 volume pp 53 - emphasis added). She then goes on to explain the nature of the prior probability: "It is very important to stress that all this (the prior probability) refers only to possibilities, none of which are yet realized at t_0 . It refers to what the decision-maker should be prepared to choose in hypothetical circumstances *before more evidence is collected*. It describes the static situation at t_0 in terms of a certain probability distribution satisfying the probability axioms." (1975, pp. 64, emphasis added).

Teller (1975 same book) argues that in order to be able to conditionalize on I_S (the relevant evidence) an initial probability function, which does not depend on I_S , is required in order for the process of updating beliefs to get started: "However, for a Bayesian theory fully to characterize the degrees of belief which arise by conditionalization, the theory must specify the belief function from which to start. This initial function is called the prior probability function" (1975, pp. 168-169). Earman (1992) also emphasizes the need of a starting point in order for the conditionalization process to be operational: "...an agent begins as a *tabula rasa*, chooses her priors, and forever after changes her probabilities only by conditionalization." (1992, pp. 139-140 original emphasis).

Cr_n for a time point T_n is uniquely determined by Cr_0 and K_n : For any H ,

$$Cr_n(H) = Cr'_0(H | K_n), \quad (4)$$

where Cr'_0 is the conditional function based on Cr_0 " (1962, pp. 310, emphasis added). Is Cr_0 actual or counterfactual? Carnap explicitly allows the initial credence function to be hypothetical: He first raises the question: "How can we understand the function Cr_0 ?" The answer to this question depends on whether the agent X is a robot or a human being. In the first case, Cr_0 is robot's actual credence function at T_0 . With respect to the second (more relevant case) Carnap's answer is as follows: "In the case of a human being X, suppose that we find at the time T_n his credence function Cr_n . Then we can, under suitable conditions, reconstruct a sequence E_1, E_2, \dots, E_n , the proposition K_n , and a function Cr_0 such that (a) E_1, E_2, \dots, E_n are possible observation data, (b) K_n is defined by (3), (c) Cr_0 satisfies all requirements of rationality for initial credence functions, (d) the application of (4) to the assumed function Cr_0 and K_n would lead to the ascertained function Cr_n . We do not assert that X *actually experienced the data* E_1, E_2, \dots, E_n , and that he *actually had the initial credence function* Cr_0 , but merely that under idealized conditions, his function Cr_n could have evolved from Cr_0 by the effect of the data E_1, E_2, \dots, E_n ." (1962, pp. 310, emphasis added)

Why does Carnap identify the initial credence function Cr_0 , instead of the current credence function Cr_n , as the basic concept upon which (any) rational decision theory must be built? The answer to this question is based on the distinction between the two competing concepts that drive a person's (say X) degrees of belief mentioned above, namely: (i) X's momentary inclination at time T and (ii) X's permanent disposition to believe. Carnap argues that what we are interested in, in developing rational decision theory is the second concept, that is a "trait" of X's "underlying permanent intellectual character". This disposition is best reflected on X's initial credence function, that is X's degrees of belief prior to the acquisition of any evidence. Any change in X's beliefs, due to the emergence of new evidence, E , should be based on $Cr_0(\cdot | E)$. Such a conditionalization ensures that X's current credences will also be driven by X's permanent disposition for forming beliefs on the basis of the received evidence. To this end, Carnap argues as follows: "When we wish to judge the morality of a person, we do not simply look at some of his acts, we study rather his character, the system of his moral values, which is part of his utility function. Single acts without knowledge of motives give little basis for a judgement. Similarly, if *we wish to judge the rationality of a person's beliefs, we should not simply look at his present beliefs. Beliefs without knowledge of the evidence out of which they arose tell us little. We must rather study the way in which the person forms his beliefs on the basis of evidence.* In other words, we should study his credibility function, not simply his present credence function." (1962, pp. 312, emphasis added).

The distinction between actual versus hypothetical priors in the context of Schmeidler's two-coin example may be described as follows: Assume that at time

t , in which the desire for a specific investigation or bet emerged (for example the time at which the agent decided to bet on D_H^A), specific information I_S already existed (for example, the coin A has already been flipped five times, with the results being H_A, H_A, T_A, H_A, T_A). What is the agent advised to do in this case? Should she form her priors by taking I_S into account? Or should she "travel back in time" and ask herself "what would have my priors been at time t_0 when I_S was not available?" MB responds negatively to the first question and positively to the second.

Based on the above MB recommendation, the following pressing question presents itself: Why should the agent take the burdensome and imagination-stretching counterfactual route (2) in forming her current beliefs, while the alternative actual (and intuitively appealing) course (1) is readily available? Put differently, why should we require the agent to ponder her subjective probabilities under an epistemic state which is not her actual state of knowledge but rather a counterfactual one? There are at least three reasons that compel an agent to follow the counterfactual path: These reasons are analyzed in detail in the next Section.

3 Arguments for Modern Bayesianism (MB)

3.1 Modern Bayesianism and the Chance-Credence Relationship: Avoiding Inconsistencies

Principal Principle (PP), one of the pillars of rationality, originally suggested by David Lewis (1970), aims at capturing the following intuitive idea: If the agent X comes to know (with certainty) that the objective probability of the event (proposition) A, $Ch(A)$, is p and does not possess any inadmissible evidence for A, then X's subjective probability of A must be set equal to p . One of the arguments of this section, is that MB is consistent with PP, whereas is not consistent with CB.

How to describe formally the PP? One option is to use the current credence function $P_t^{I_B, I_S}$, which gives:

$$P_t^{I_B, I_S}(A | \langle Ch(A) = p \rangle, E) = p, \quad (5)$$

where E is admissible evidence obtained at $t + 1$.

It must be noted that Lewis never gave a precise definition of "admissibility". The standard interpretation of this concept is "that evidence is admissible if it is not relevant to the outcome of the chance event in question. As a rule of thumb, he (Lewis) takes information about the past to be admissible, and information about the future to not be admissible." (Meacham 2007, pp. 18).

The first version of PP, given by (5), may be interpreted in the light of CB, as follows: The agent X, being at time t , has at her disposal I_B and I_S . Being a CB, the agent uses both these items of information to generate her "unconditional" probability $P_t^{I_B, I_S}(A)$ at t . If she acquires a further item of

information E then her new updated probability of A would be

$$P_{t+1, I_B, I_S, E}(A) = P_t^{I_B, I_S}(A | E).$$

However, if in addition to E the agent learned that the chance of A is p , then the latter piece of information would dominate the formation of her beliefs, in the sense that $\langle Ch(A) = p \rangle$ screens-off any admissible evidence, E , for A, thus yielding (5). Hence, her new (updated) probability of A is given by,

$$P_{t+1, I_B, I_S, E, \langle Ch(A)=p \rangle}(A) = P_t^{I_B, I_S}(A | \langle Ch(A) = p \rangle, E) = p. \quad (6)$$

A second option for the agent is to operate as an MB, thus treating I_S not as actual but rather as counterfactual. As a result, she employs the initial prior probability function $P_0^{I_B}$, instead of the current one, thus yielding the following version of PP:

$$P_0^{I_B}(A | \langle Ch(A) = p \rangle, E, I_s) = p \text{ if } E \text{ (and } I_s \text{) is admissible.} \quad (7)$$

What is the restriction that (7) imposes on X's current credences? Put differently, assume X is currently at time $t + 1$, in which she has come to know I_S , $\langle Ch(A) = p \rangle$ and E . MB's condition (2) implies the following relation:

$$P_{t+1, I_B, I_S, E, \langle Ch(A)=p \rangle}(A) = P_0^{I_B}(A | \langle Ch(A) = p \rangle, E, I_s) = p. \quad (8)$$

Comparing (6) with (8), we form the impression that it makes no difference to the agent's updated probability of A, whether she acts as a CB or MB. However, Meacham (2007) argues that this is not the case. He argues that (5), as opposed to (7), leads to contradictions. To understand why, we must first emphasize the difference that the role that I_S plays in (5) versus (7): On one hand, under the MB formulation (7) of PP, I_S is conditioning information and therefore, we must examine if it is admissible or not. Now, let us consider the special case in which A is included in I_S at time t . In such a case, I_S becomes automatically inadmissible and therefore, PP becomes non-applicable and the agent's probability should be set equal to one.

On the other hand, I_S is not conditioning information; it is rather a direct determinant of the agent's current probability function $P_t^{I_B, I_S}$. As such, admissibility is not relevant. Therefore, the agent cannot condition on A since it has already determined her probability function. Hence, (5) dictates that the agent's subjective probability of A should be equal to $p < 1$. But on the other hand, this probability should be equal to one, since A is known to the agent. Hence, we have run into the following contradiction:

$$1 = P_t^{I_B, I_S}(A | \langle Ch(A) = p \rangle, E) = p \neq 1$$

Obviously the extent to which the prior probability is preferable to the current probability for both formal and conceptual reasons, determines the extent to which MB is preferable to CB.

3.2 Modern Bayesianism and the Problem of Old Evidence

The "problem of old evidence" concerns the role of evidence, E , that was already known (hence, old) at the time of the formation of a scientific theory (T) for the confirmation (or not) of T . The problem was first presented by Glymour (1980) who described it as follows: "Scientists commonly argue for their theories from evidence known long before the theories were introduced. Copernicus argued for his theory using observations made over the course of millennia.... Newton argued for universal gravitation using Kepler's second and third laws, established before the Principia was published. The argument that Einstein gave in 1915 for his gravitational field equations was that they explained the anomalous advance of the perihelion of Mercury, established more than half a century earlier.... Old evidence can in fact confirm new theory, but according to Bayesian kinematics it cannot. For let us suppose that evidence e (E in our notation) is known before theory T is introduced at time t . Because e is known at t , $P_t(e) = 1$. Further, because $P_t(e) = 1$, the likelihood of e given T , $P_t(e|T)$, is also 1. We then have:

$$P_t(T|e) = \frac{P_t(T)P_t(e|T)}{P_t(e)} = P_t(T)$$

The conditional probability of T on e is therefore the same as the prior probability of T : e cannot constitute evidence for T None of the Bayesian mechanisms apply, and if we are strictly limited to them, we have the absurdity that old evidence cannot confirm a new theory." (1980, pp. 85-86). The previous quotation highlights a conflict between Bayesian theory of confirmation and the standard scientific practice when it comes to the use of E . Specifically, E , under the standard scientific practice, can in fact confirm new theory, but according to Bayesian theory it cannot.

Bayesian philosophers have responded to the problem of old evidence in (mainly) two different (and incompatible) ways. The first response, put forward by Garber (1983), Niiniluoto (1983) and Jeffrey (1983) goes as follows: What makes the scientist to increase her confidence in T , is not E per se, since this information was already known to her at the time she invented her theory. Rather it was the realization of the *logical entailment relationship*

$$T \vdash E \tag{9}$$

that made the scientist to increase her confidence in T . For example, "Einstein discovered only after writing down the equations of General Relativity that they entailed the anomalous perihelion advance of Mercury." (Howson, 1991, pp. 553). Hence, the increase in the agent's degree of belief in T comes from conditionalization on the entailment relationship (9) itself rather than conditionalization on just E . The proposition (9) at t_0 was uncertain, in the sense that the agent had not established/proved the validity of the postulated entailment relationship, i.e.

$$P_0(T \vdash E) < 1,$$

in which case the realization of (9) at t results in an increase in probability of T .

Note that this line of defense is based on the assumption that the agent conditionalizes on the entailment (9), which in turn implies that at t_0 , the agent had already assigned a prior probability to this proposition, i.e. $P_0(T \vdash E) = p$, with $p < 1$. This means that at t_0 the agent was not aware of the truth of (9) (which truth she discovered only later) which in turn implies that at t_0 the agent had assign a probability less than one to a logical truth. This, however, means that P_0 is not coherent. Hence, this line of defense against the problem of old evidence leads to another problem (potentially more serious) which is incoherence of the agent's probability function.

The second argument for solving the problem of old evidence is more in the spirit of MB since it is based on the following counterfactual "what-if" strategy. Garber (1983) portrays this strategy as follows: "One obvious response might begin with the observation that *if one had not known the evidence in question*, then its discovery *would have increased one's degrees of belief* in the hypothesis in question (H). That is, in the circumstances in which E really does confirm H , if it had been the case that $P(E) < 1$, then it would also have been the case that $P(H|E) > P(H)$." (1983, pp. 103, emphasis added). The solution outlined by Garber, suggests that the agent should not evaluate her probabilities relative to $I = I_B \cup I_S$ but only relative to I_B . This is exactly the proposal offered by Howson (1991) in his attempt to solve the problem of old evidence. Indeed, he explicitly states that the probabilities "should always be relativized to $K - \{E\}$ " (1991, pp. 548, with K and E in Howson's notation corresponding to I_B and I_S , respectively). But, as analyzed in Introduction, this is exactly the point at which MB and CB differ. Hence, choosing MB over CB might be motivated by the fact that the first strategy fares better than the second in cases of "old evidence".

3.3 Modern Bayesianism and Psychological Detachment from Actual Evidence

Let us now turn our attention to the third reason for preferring the counterfactual MB strategy over the actual CB one. This reason is akin to the possibility that the agent evaluates the entailment relationships between theoretical and evidential statements in a more unbiased way when the evidence I_S is hypothetical than when it is actual. For example, the probability that I assign to the hypothesis H_A : "I shall live for at least another five years" conditional on the actual event B : "I am diagnosed with aggressive lung cancer" is likely to be bigger than the probability that I would have assigned to H_A if the event B were not actual but just a member of a set of many hypothetical events including the event non- B . In this example, the actual knowledge of the unfortunate event B at time t (now) is likely to create psychological pressure to distort the partial entailment relationship between H_A and B that I was willing to accept at t_0 , thus yielding

$$P_{B,t}(H_A) > P_0(H_A | B).$$

This argument suggests that the subjective evaluation of probabilities should be made, as Longino (1979, pp. 35) puts it, "in rational, cool, moments". As already mentioned, such a "cool" and "impartial" agent is an MB agent.

4 Objections to Modern Bayesianism: Logical Omniscience

The arguments presented above suggest that MB is the optimal strategy that must be followed by an agent who is in the process of forming her subjective probability function at t (now). For reasons of clarity, this strategy is summarized below. The agent is currently at time t , knowing both I_B and I_S . The agent is advised to "pretend" that she does not know I_S and perform the following thought experiment: Go back in time at period t_0 in which only I_B was available. At t_0 , the background information I_B contains the full set of hypotheses \mathbf{H} that the agent actually entertained at t_0 , the full set of evidential sentences \mathbf{E} that the agent conceived as possible at t_0 , (every conceivable possible course events after t_0) as well as all the entailment relationships between elements of \mathbf{H} and elements of \mathbf{E} (e.g. $H_i \vdash E_j$). More specifically, each member of \mathbf{E} describes the observations that have not been actually made at t_0 but are considered by the agent at t_0 as possible to make at the future period t . At t_0 the agent's subjective probability function P_0 was defined over the field of propositions \mathcal{L} (which includes \mathbf{H} and \mathbf{E}). Assume that at t , the particular element $E \in \mathbf{E}$ is realized, which implies that the agent at t has this specific information, i.e. $I_S \equiv E$. According to the MB counterfactual strategy, the agent's new probability function, P_t , at t should be her old probability function P_0 conditional on E . As a result, her unconditional new probability for the uncertain event A (that will occur at $t+1$) should be set equal to the conditional probability $P_0(A | E)$ rather than being determined "from scratch" as $P_{t,E}(A)$.

Is the above counterfactual strategy without any problems? The answer is negative. The main objections raised against this strategy revolve around the concept of logical omniscience (or its lack thereof). Specifically, observe that in the description of the counterfactual strategy presented above, we defined \mathbf{H} and \mathbf{E} as the sets of theoretical and evidential statements, respectively *conceived by the agent* at t_0 . However, Bayesianism assumes something significantly more than that. It assumes that the agent knows *all* the possible hypotheses, \mathbf{H}^* , *all* the possible evidential statements \mathbf{E}^* and *all* the possible entailment relationships between \mathbf{H}^* and \mathbf{E}^* ; in other words the agent is assumed to be logically omniscient and her prior probability function, P_0^* , is defined over the field \mathcal{L}^* of *all propositions*, which is (infinitely) richer than \mathcal{L} . Why does Bayesianism require such an extreme condition? The answer is to ensure coherence of P_0^* . Specifically, if we wish (as Bayesianist do) to ensure that P_0^* is a proper probability measure then we must establish that to each logical consequence $p \vdash q$

the agent assigns probability equal to one, that is

$$\forall(p \vdash q), P_0^*(p \vdash q) = 1.$$

This is because if p logically entails q , it does so in every possible state of the world. Hence, it appears that coherence requires that the agent is capable of tracking all logical consequences at t_0 . In other words, the agent is not allowed to be unaware even of a single logical truth that exists within the underlying domain \mathcal{L} . As will be shown below only under this extreme form of logical omniscience the aforementioned MB counterfactual strategy is expected to work under all epistemic states that the agent may entertain at t .

4.1 Local vs. Global Bayesianism

Logical omniscience of the form described above is a condition that is usually met within the variant of Bayesianism, known as Global Bayesianism (GB). Garber (1983) defined GB as follows: "One popular conception of the Bayesian enterprise is what I shall call global Bayesianism. On this conception, what the Bayesian is trying to do is build a global learning machine, a scientific robot that will digest all of the information we feed it and churn out appropriate degrees of belief. On this model, the choice of a language (domain) over which to define one's probability function is as important as the constraints that one imposes on that function and its evolution. On this model, *the appropriate language to building into the scientific robot is the ideal language of science, a maximally fine grained language \mathcal{L} , capable of expressing all possible hypotheses, all possible evidence, capable of doing logic, mathematics, etc.* In short, \mathcal{L} must be capable, in principle, of saying anything we might ever find a need to say in science." (1983, pp. 110 emphasis added).

In what respect, does the MB strategy suffer in the case of an agent whose prior probability function is P_0 rather than P_0^* ? This question is equivalent to the following one: "Why is GB not realistic?". The problem is the following: Assume that at t the agent comes across a piece of evidence E^* that she had not conceived of as possible at t_0 . Assume that this piece of evidence is the agent's specific information (I_s) obtained at t . According to MB, the agent should go back to t_0 and conditionalize on E^* in order to generate her new probability function at t . However, an unpleasant surprise awaits the agent: At t_0 there is no probabilistic assignment on E^* simply because E^* was not in the domain of P_0 at t_0 . To make things even worse, assume that E^* was obtained at t because a new hypothesis, H^* , was put forward at t which motivated the observations (or experiments) that led to E^* . Obviously, H^* was not in the domain of P_0 either. For example, consider as E^* and H^* the "deflection of light by the sun" and "Einstein's general relativity theory", respectively. If a physicist of the late nineteenth century was asked to assign his personal probability to the event that the light is deflected by the sun, he would have been caught by surprise. His domain, \mathcal{L} at t_0 was simply not rich enough to accommodate either E^* or H^* (let alone the entailment $H^* \vdash E^*$). In such a case, counterfactual conditionalization seems impossible and the whole MB strategy runs into trouble.

The preceding analysis suggests that the MB counterfactual strategy is unrealistic to the same extent as GB, (as a description of the epistemic state of real people) is. If the MB strategy is unrealistic within the GB framework, then is it possible that there is another framework, say it Local Bayesianism (LB), within which the MB strategy is likely to work? Before we answer this question, let us first describe, following Garber (1983) how LB might be defined: "Typically when scientists or decision makers apply Bayesian methods to the clarification of inferential problems, they do so in a much more restricted scope than global Bayesianism suggests, *dealing only with the sentences and degrees of belief that they are actually concerned with, those that pertain to the problem at hand.* This suggests a different way of thinking about the Bayesian learning model, what one might call *local Bayesianism.* On this model, the Bayesian does not see himself as trying to build a global learning machine, or a scientific robot. Rather, the goal is to build a hand-held calculator, as it were, a tool to help the scientist or decision maker with particular inferential problems. On this view, the Bayesian framework provides a general formal structure in which one can set up a wide variety of different inferential problems. In order to apply it in some particular situation, *we enter in only what we need to deal with in the context of the problem at hand, i.e., the particular sentences with which we are concerned, and the beliefs (prior probabilities) we have with respect to those sentences.*" (1983, pp. 111, emphasis added).

Formally, assume that the investigator is interested in the set of specific hypotheses (theoretical propositions) H_i , $i = 1, 2, \dots, n$ that cover the full set of possibilities for the problem at hand (e.g. the coin is fair, the coin is biased towards H, the coin is biased towards T, etc). Related to the aforementioned theoretical propositions are the set of evidential propositions E_j , $j = 1, 2, \dots, m$ which are relevant for H_i (e.g. the outcome of the next toss is H, the outcomes of the next two tosses are H,T etc). No other proposition enters the agent's "problem over relative domain" \mathcal{L} that is relevant for the specific problem at hand. More specifically, Garber (1983, pp. 111) defines \mathcal{L} to be "just the truth-functional closure" of the H_i , E_j and propositions of the logical entailment between H_i and E_j , that is propositions of the form " $H_i \vdash E_j$ ". As a result, the agent's prior probability function is defined over the "local" more modest domain \mathcal{L} rather than the "global" ideal domain \mathcal{L}^* . This in turn implies that it is much easier for the agent to track all logical consequences " $H_i \vdash E_j$ " (for some i and j) within \mathcal{L} than within \mathcal{L}^* , thus making local Bayesianism more realistic than global Bayesianism (Bayesianism with a human face, as Jeffrey (1983) puts it).

4.2 Small vs. Large Worlds

The description presented above makes clear that LB refers to well-defined, small-scale cases (characterized by a small number of hypotheses, evidential propositions and logical relationships) which are usually referred to as "small worlds". In such a world, it is quite plausible to assume that the agent can consider all the possibilities (hypotheses, evidential propositions and their en-

tailment relations) that are relevant for the problem at hand at t_0 . In such a case, no surprises await the agent at t , in the sense that his epistemic state at t is identical to that at t_0 . Binmore (2009) defines a small world as one "within which all potential surprises have been predicted and evaluated in advance of their occurrence." (2009, pp 8).

What are the implications of small worlds for the effectiveness of the MB counterfactual strategy? Binmore (2009) comments on this question as follows: "Only in a small world, in which you can always *look before you leap*, it is possible to consider everything that might be relevant to the decisions you take." (2009, pp. 139 emphasis added). Indeed the "look before you leap" proverb is attributed to Savage who used it as antithetical to "cross that bridge when you come to it" that referred to the so called "large worlds". It is worth mentioning that Savage himself made quite clear that his own conception of subjective probability together with its axiomatization is relevant only for small worlds. This is because Savage's framework is essentially static in the sense that it does not allow for the so-called "concept formation", that is the formation of a new hypothesis or a new idea sometime in the future. Commenting on Savage's approach, Suppes (1966) observes: "The important thing I wish to emphasize is that the theory provides no place for the decision-maker to acquire a new concept on the basis of new information received. The theory is static in the sense that it is assumed that the decision-maker has a fixed conceptual apparatus available to him throughout time." (1966, pp. 21). Savage himself acknowledged the fact that the static nature of his theory makes it inapplicable in the case of large evolving worlds by referring to such an extension as "ridiculous" and "preposterous".

5 Small Worlds, Counterfactual Strategy and Ambiguity Aversion

The aforementioned discussion may be summarized as follows: MB and CB refer to the agent's strategy or course of action in forming her probability function at t . GB and LB refer to the agent's epistemic status at t_0 as reflected in the contents of her background information I_B or her domain, (\mathcal{L} or \mathcal{L}^*). Whether MB or CB is appropriate depends on which epistemic status, GB or LB characterizes the agent. Obviously there are four combinations for the agent to consider in her attempts to solve the problem of forming her subjective probability at t' :

- (i) MB-GB: Ideal but Unrealistic (Infeasible) Solution.
- (ii) MB-LB: The Appropriate Solution in Small Worlds.
- (iii) CB-GB: A Feasible Solution in Large Worlds but with the kind of problems analyzed in section 3.
- (iv) CB-LB: Inappropriate or Inefficient Solution.

How does the above classification bear upon AA? The preceding discussion has made clear that the fact that the agent allowed I_S to affect directly her probabilistic beliefs, is only a necessary condition for AA. A further condition is

that the problem at hand falls into the category of large worlds. Let us clarify this argument, which forms our basic thesis in this paper. Assume that the agent at time t has information, I_S , about the objective probabilities of the events in \mathcal{F}_k , but she lacks similar information about the events in \mathcal{F}'_k . If the world is small, the agent can go back in time at t_0 , ignore I_S and assign probabilities in $\mathcal{F}_k \cup \mathcal{F}'_k$ in an information-symmetrical way. Once her prior probability is thus determined (counterfactually at t_0), the agent is allowed to bring I_S back to the picture by conditionalizing on it. In such a case, the only way for the agent to produce AA is to deviate at t from her probabilistic commitments at t_0 , that is by exhibiting dynamic inconsistency. Interestingly, all the cases motivating AA that have appeared in the literature (including the original Ellsber's paradox as well as Schmeidler's two coin example) are "textbook cases" of small worlds. Hence, although we do not claim that AA cannot arise in large worlds, (because in such a context the counterfactual MB strategy is unrealistic) we do argue that AA has been poorly motivated!

6 Schmeidler's two-coin example Revisited

Let us now revisit Schmeidler's two-coin example discussed in the introduction in light of the analysis presented in section 4. First describe the agent's epistemic background at t_0 . For simplicity, we assume that concerning coin A, the agent knows with certainty that only one of the following three hypotheses is true:

$$\mathcal{H}_1^A = \{\text{Coin A is fair}\}$$

$$\mathcal{H}_2^A = \{\text{Coin A favors H}\}$$

$$\mathcal{H}_3^A = \{\text{Coin A favors T}\}.$$

To make things even simpler assume that in the case of \mathcal{H}_2^A the objective probability of H is equal to 0.6, whereas in the case of \mathcal{H}_3^A the objective probability of T is 0.6. Similar assumptions are made for coin B, namely the agent knows that one of the following three hypotheses is true:

$$\mathcal{H}_1^B = \{\text{Coin B is fair}\}$$

$$\mathcal{H}_2^B = \{\text{Coin B favors H}\}$$

$$\mathcal{H}_3^B = \{\text{Coin B favors T}\}$$

Let $\mathbf{H}_A = \{\mathcal{H}_1^A, \mathcal{H}_2^A, \mathcal{H}_3^A\}$ and $\mathbf{H}_B = \{\mathcal{H}_1^B, \mathcal{H}_2^B, \mathcal{H}_3^B\}$.

The agent has to decide about her prior probabilities of the above hypotheses. Since there is no direct information supporting one hypothesis over the rest, the

agent is likely to subscribe to the principle of indifference (or the maximum entropy principle) according to which:

$$P_0(\mathcal{H}_1^A) = P_0(\mathcal{H}_2^A) = P_0(\mathcal{H}_3^A) = \frac{1}{3}$$

and

$$P_0(\mathcal{H}_1^B) = P_0(\mathcal{H}_2^B) = P_0(\mathcal{H}_3^B) = \frac{1}{3}.$$

Alternatively, the agent might consult her past experience on similar cases (part of her background information) which suggests that it is more often for someone to encounter cases of fair coins than not. In any case, the important thing to notice is that there is no specific information at t_0 , hence there is no informational asymmetry between the set of hypotheses concerning coin A and that of coin B.

Let us now calculate the probabilities that the agent assigns to the events/propositions D_H^A , D_T^A and D_H^B , D_T^B defined as:

$$\begin{aligned} D_H^A &= \{\text{coin A comes up heads}\} = \{H_A H_B, H_A T_B\} \\ D_T^A &= \{\text{coin A comes up tails}\} = \{T_A H_B, T_A T_B\} \\ D_H^B &= \{\text{coin B comes up heads}\} = \{H_A H_B, T_A H_B\} \end{aligned}$$

and

$$D_T^B = \{\text{coin B comes up tails}\} = \{H_A T_B, T_A T_B\}$$

Using the law of total probability, the agent gets:

$$P_0(D_H^A) = \sum_{i=1}^3 P_0(D_H^A | \mathcal{H}_i^A) P_0(\mathcal{H}_i^A) = 0.5 \times \frac{1}{3} + 0.6 \times \frac{1}{3} + 0.4 \times \frac{1}{3} = 0.5.$$

In a similar fashion, we obtain,

$$P_0(D_T^A) = P_0(D_H^B) = P_0(D_T^B) = 0.5.$$

Let us assume that at time $t > t_0$ the agent acquires an important piece of specific information for the problem at hand. In particular she is given a set of data, E_∞^A , consisting of the results of a very long series of tosses of coin A for which (almost) half of them are H. The following two questions are of interest: (i) how does the agent update her probabilities of D_H^A , D_T^A , D_H^B and D_T^B in light of the new information, E_∞^A about A and (ii) does she end up with a new system of beliefs that violate coherence?

Bayesian conditionalization implies:

$$P_t(D_H^A) = P_0(D_H^A | E_\infty^A) = \sum_{i=1}^3 P_0(D_H^A | \mathcal{H}_i^A, E_\infty^A) P_0(\mathcal{H}_i^A | E_\infty^A)$$

The Principal Principle implies that (knowledge of the objective chance screens off any admissible evidence):

$$P_0(D_H^A \mid \mathcal{H}_i^A, E_\infty^A) = P_0(D_H^A \mid \mathcal{H}_i^A), \text{ for } i = 1, 2, 3. \quad (10)$$

Furthermore, given that the evidence E_∞^A is assumed to be arbitrarily large:

$$\begin{aligned} P_0(\mathcal{H}_1^A \mid E_\infty^A) &\simeq 1 \\ P_0(\mathcal{H}_2^A \mid E_\infty^A) &\simeq 0 \\ P_0(\mathcal{H}_3^A \mid E_\infty^A) &\simeq 0 \end{aligned}$$

The above relationships together with (10) imply that

$$P_t(D_H^A) \simeq P_0(D_H^A \mid \mathcal{H}_1^A)P_0(\mathcal{H}_1^A \mid E_\infty^A) = 0.5$$

Similarly,

$$P_t(D_T^A) = 0.5.$$

Let us now turn our attention to the events, D_H^B and D_T^B (related to coin B) and examine how the acquisition of E_∞^A affects their probabilities.

$$P_t(D_H^B) = P_0(D_H^B \mid E_\infty^A) = \sum_{i=1}^3 P_0(D_H^B \mid \mathcal{H}_i^B, E_\infty^A)P_0(\mathcal{H}_i^B \mid E_\infty^A)$$

As before, PP implies

$$P_0(D_H^B \mid \mathcal{H}_i^B, E_\infty^A) = P_0(D_H^B \mid \mathcal{H}_i^B), \text{ for } i = 1, 2, 3.$$

Now the crucial question is the following: Does the evidence on coin A affect the prior probabilities attached to coin-B related events? Assuming independence between the sets \mathbf{H}_A and \mathbf{H}_B the answer is negative. Hence,

$$\begin{aligned} P_0(\mathcal{H}_1^B \mid E_\infty^A) &= P_0(\mathcal{H}_1^B) = \frac{1}{3} \\ P_0(\mathcal{H}_2^B \mid E_\infty^A) &= P_0(\mathcal{H}_2^B) = \frac{1}{3} \\ P_0(\mathcal{H}_3^B \mid E_\infty^A) &= P_0(\mathcal{H}_3^B) = \frac{1}{3} \end{aligned}$$

Finally, the above relationships imply

$$P_t(D_H^B) = P_t(D_T^B) = 0.5.$$

The important thing to notice is that the agent's new set of beliefs, P_t , obtained at t (after the specific information E_∞^A has arrived) is coherent, that is, it is a proper probability function, i.e. it does not violate the additivity property.

Why did the preceding analysis not produce AA? Put differently, why the acquisition of E_∞^A did not produce a system of beliefs having the property

$$P_t(D_H^B) < P_t(D_H^A)$$

and

$$P_t(D_T^B) < P_t(D_T^A)$$

which would imply an ambiguity averse, probabilistically non-sophisticated agent? The answer is that the analysis presented above implies that the agent follows the MB counterfactual strategy in forming her probability function P_t at t . Indeed, the key point is to focus on how the agent formed her new probability function P_t once she became aware of E_∞^A . Specifically, the agent did not use E_∞^A as a direct determinant (on a par with background information) of P_t . Instead, she "traveled back in time" at point t_0 , identified P_0 and stuck to those commitments at t^4 . In this setting, AA would emerge only if the agent abandoned her commitments made at t_0 and assigned different probabilities at t , thus exhibiting dynamic inconsistency.

The preceding analysis implies that the only way for AA to emerge (without causing dynamic inconsistency) is to eliminate P_0 . Eliminating P_0 means that at t the agent realizes that she faces a new reality that had not (and could not have) anticipated at t , which obliges her to re-evaluate her full set of probabilistic beliefs (to form a new prior). Such a setting can arise within a large world framework. Teller (1973) gives an example of when such a strategy is required: "Examples are the wildcatter's problem of where to drill an oil well and the manufacturer's problem of how much to produce. The personalists advise the wildcatter and the manufacturer to estimate their initial degrees of belief subjectively for any one such problem, *but if a similar problem arises years later under considerably different circumstances they are advised to make new evaluations rather than to try to conditionalize their old estimates on the vast body of intervening observations of uncertain relevance to the problem.*" (1975, pp. 170, emphasis added). However, Schmeidler's and Ellsberg's examples analyzed above are quite different from the wildcatter's problem. Consider the following question with respect to Schmeidler's two coin example: Is there any reason for the agent to refine her probabilistic judgements at t , because at t she came across a piece of surprising information, that is information that she could not have conceived back at t_0 ? The answer is negative, because Schmeidler's two coin example takes place in a small world, where the concept of "surprising information" is not relevant. As a result, the MB strategy is perfectly feasible; in fact it follows quite naturally!

7 Concluding Remarks

We conclude by summarizing our main arguments. There are two ways to use specific information, I_S , in constructing the agent's subjective distribution at t , P_t . The first is to use it as conditioning information, consistent with MB while the second is to use it on a par with any other background information I_B , consistent with CB. There are good reasons, analyzed in detail in Section

⁴Note that the agent having formed P_0 at t_0 , was at an epistemic state of not knowing with certainty that E_A^∞ has occurred

3, for preferring the first interpretation (MB) over the second (CB). Does the choice of MB over CB have implications for the main issue under scrutiny in the present paper, namely the emergence of AA? The answer is yes. If we choose MB there is no room for AA, since being ambiguity averse in such a framework is equivalent to being dynamically inconsistent. On the other hand, AA can arise under CB because when I_S determines the agent's subjective distribution, leads to informational asymmetries.

Can we always choose MB over CB? The answer is negative. MB is applicable only in small worlds. In large worlds, MB may not be operational, especially when the agent's epistemic state at the time that I_S arrives is richer than the one at the time prior to the arrival of I_S , which is the case in large worlds. Hence, AA may be thought of as a (potential) property of large worlds. To this end, the motivation of AA in the form of examples such as Ellsberg's classic urn-based or Schmeidler's coin-based, is clearly poor since these examples are textbook cases of small worlds.

References

- Bammer, G. and M. Smithson (2008). "Uncertainty and Risk: Multidisciplinary Perspectives". London: Earthscan.
- Binmore, K. (2009). "Rational Decisions, Princeton". Princeton University Press.
- Carnap, R. (1950). "Empiricism, Semantics and Ontology", *Revue Internationale de Philosophie* 4, 20-40. Reprinted in *Meaning and Necessity: A Study in Semantics and Modal Logic*, 2nd edn. University of Chicago Press, Chicago.
- Carnap, R. (1962). "Logical Foundations of Probability". 2nd ed. Chicago: University of Chicago Press. Par. 87, 468-478.
- Easwaran, K. (2008). "The Foundations of Conditional Probability". PhD thesis, University of California, Berkeley.
- Ellsberg, D. (1961). "Risk, ambiguity, and the Savage axioms", *Quarterly Journal of Economics* 75, 643-669.
- Garber, D. (1983). "Old Evidence and Logical Omniscience in Bayesian Confirmation Theory", in J. Earman (ed): *Testing Scientific Theories* 99-132. Minnesota Studies in the Philosophy of Science, vol. X. Minneapolis: University of Minnesota Press.
- Gilboa, I. and D. Schmeidler (1989). "Maxmin expected utility with a nonunique prior", *Journal of Mathematical Economics* 18, 141-153.
- Glymour, C. (1980). "Theory and Evidence". Princeton: Princeton University Press.
- Howson, C. (1991). "The 'old evidence' problem", *British Journal for the Philosophy of Science* 42: 547-55.
- Howson, C. and P. Urbach (1991). "Bayesian reasoning in science", *Nature* 350:371-374.
- Jeffrey, R. C. (1983). "The Logic of Decision". 2nd ed. Chicago: University of Chicago Press.
- Lewis, D. (1973). "Counterfactuals". Oxford: Blackwell.
- Longino, H. (1979). "Evidence and hypothesis", *Philosophy of Science*, 46 (1): 35-56.
- Machina, M. J. and S. Marciiano (2014). "Ambiguity and Ambiguity Aversion". *Handbook of the Economics of Risk and Uncertainty*, 729-807, Elsevier.
- Meacham, J. (2007). "American Gospel: God, the Founding Fathers, and the making of a nation". New York, NY: Random House.

- Meacham, J. (2008). "Sleeping Beauty and the Dynamics of De Se Beliefs". *Philosophical Studies* 138: 245–70.
- Niiniluoto, I. (1983). "Novel Facts and Bayesianism", *The British Journal for the Philosophy of Science* 34: 375–379.
- Savage, L. (1954). "The foundations of statistics". New-York, John Wiley.
- Schmeidler, D. (1989). "Subjective probability and expected utility without additivity", *Econometrica* 57(3), 571–587.
- Suppes, P. (1966). "Some formal models of the grading principle", *Synthese* 6, 284–306.
- Viscusi, W. K. and H. Chesson. (1999). "Hopes and Fears: The Conflicting Effects of Risk Ambiguity", *Journal of Risk and Uncertainty* 47, 153–178.