DEPARTMENT OF INTERNATIONAL AND EUROPEAN ECONOMIC STUDIES

ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

Amgiguity Aversion, Modern Bayesianism and Small Worlds

Phoebe Koundouri

Nikitas Pittis

Panagiotis Samartzis

Nikolaos Englezos

Andreas Papandreou

Working Paper Series

17-02

July 2017

Amgiguity Aversion, Modern Bayesianism and Small Worlds

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

Abstract

1 Introduction

Since 1961, year at which Ellsberg introduced the concept of "ambiguity aversion" (AA) in the form of his well known thought experiments, economists, psychologists and decision theorists have been trying to provide answers to the following two questions: (i) is in fact AA empirically documented by means of formal experimentation and (ii) is ambiguity aversion a "paradox" in the sense that it violates the standard Bayesian conditions of rationality, or is it the other way round, namely AA is consistent with rationality provided that the latter is modified to make room for such a behavior?

With respect to the first question, Machina and Siniscalchi (2013) offer an extensive survey of the empirical literature (including studies from the insurance and medical literature) that spans a period of about 50 years. The results of these experiments seem to lend support to AA, although some results that point towards "ambiguity neutral" or even "ambiguity seeking" behaviour do exist (see, for example, Viscusi and Chesson 1999).

*School of Economics, Athens University of Economics and Business, Greece; Grantham Research Institute, London Schoopl 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 the theory that AA is consistent with rational probabilistic beliefs may be stated as follows: 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 theory 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 of information: Consider an event space $\mathcal{F}$ which contains the "events of interest" for the agent S. Assume that $\mathcal{F}_k \subset \mathcal{F}$ contains those events whose probabilities are known to S, 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 S has no information about their probabilities. AA means that S favors betting on events in $\mathcal{F}_k$ than $\mathcal{F}'_k$. However, under the maintained assumption that S's subjective probability of any event $D \in \mathcal{F}$ is measured by S's willingness to bet on D, AA implies that S's degrees of belief defined on $\mathcal{F}$ violate the additivity axiom of probability theory, thus falling short of being a proper probability function.

What is the main reason for the absence of the additivity property in S's degrees of belief, denoted by P_S? The answer is the presence of information about the chances of events in $\mathcal{F}_k$ and the lack of similar information for the events in $\mathcal{F}'_k$. Put differently, the argument that leads to the non-addiditivity property of P_S is based on the premise that the agent posseses *specific* (rather than only background) information, which is relevant for some specific events in $\mathcal{F}$ and irrelevant for all the rest. However, if we assume that such specific information, I_S, is available, then the following question arises: does P_S represent the *prior* beliefs of the agent? Put differently, is specific information 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 200?) 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 (more on this below). On the other hand, I_B should not include any actual outcomes of the experiment or information about the objective chances (e.g. coin j 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 the following: 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. Sometimes this debate is expresed in an equivalent form: The advocates of Option II argue that the agent's current probabilities (credences), P_t, are elligible to play the role of the "prior" for the evidence that will emerge at $t+1$. As Howson and Urbach (1991) put it "the priors in any calculation can also be posterior probabilities calculated from the historical data" (1991, pp. 374). As such, current credences 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 in the sense that, for every (future) $A \in \mathcal{F}$, $P_t(A \mid I_S) = P_t(A)$. Notationally, we may emphasize this, by replacing P_t with $P_t^{I_B,I_S}$. Following Meecham (2007), we shall refer to this view as "Classical Bayesianism" (CB).

On the contrary, the advocates of the opposite view, hereafter refered to as "Modern Bayesianism" (MB), insist on the prior, P_0, being determined at a time in which no specific evidence I_S was available. 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 was available. Hence, the prior and current probability functions differ for every $t \neq t_0$. Indeed, at t_0, I_S should be treated as hypothetical evidence, being on a par with any other evidential sentence that the agent might have envisaged at t_0. In other words, the formation of the agent's prior must be independent of the time in which the agent is actually located at the beginning of her investigations. Instead, it must be generated at a time in which all specific evidence for the events of interest is contigent rather than actual. Notationally, P_0 may be replaced by $P_0^{I_B}$.

Does the distinction between Classical and Modern Bayesianism bear any implications for the topic of our current investigation, namely whether or not AA may be made consistent with rationality? The answer to this question is emphatically affirmative. The main point of the present paper is to show that for an MB agent, operating in the context of a "small world" (more on this below), AA is impossible to arise. Put differently, the assumption that the agent is capable of "counterfactual probabilistic reasoning" (the main assumption of MB) entails the absence of ambiguity aversion. Hence, the important question seems to be, not whether the agent exhibits AA, but rather, whether the agent can (descriptive aspect) or should be able to (normative aspect) contemplate her degrees of beliefs under an epistemic state which in fact she is not in.

Is MB a normatively appealing hypothesis? To this end, we put forward three arguments in favor of MB, analyzed in detail in the following section. In a nutshell, an MB agent (as oposed to a CB one) (i) reveals her permanent dispositions rather than momentary inclinations in forming beliefs, (ii) protects herself from falling into logical contradictions akin to the relation between her subjective probability and objective chance, and (iii) ...

What about MB's descriptive status? Is MB reasonably realistic as a model

of actual behaviour? 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 language, $\mathcal{L}_t$, is richer than her language $\mathcal{L}_0$ at the beginning of her investigations (that is before any specific information I_S was available). This may happen if between the time periods 0 and t, a new hypothesis H_{new} was conceived by the agent. In such a case, H_{new} is part of $\mathcal{L}_t$ 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 $\mathcal{L}_0$. This is bad news for MB. Specifically, MB requires the agent to contemplate her probabilistic assignments (in the case that I_S were not known) not in the "old" language $\mathcal{L}_0$ (which did not include H_{new}) but in the "new", expanded language $\mathcal{L}_t$. The shift from $\mathcal{L}_0$ to $\mathcal{L}_t$, however bears an unpleaseant consequence: the agent is forced to change all her probabilistic assignments made in the context of $\mathcal{L}_0$. As Hajek (200?, pp ?? - A Philosopher's guide to probability) remarks "...changing the language in which items of evidence and hypotheses are expressed will typically change the confirmation relations between them." This is the point at which MB loses its normative and descriptive edge. Does MB's loses translate to CB's gains? The answer is negative. Both MB and CB are equally hurt by the presence of large and evolving worlds since in both contexts, frequent changes in the prior probabilities $P_0^{I_B}$ and $P_t^{I_B,I_S}$, respectively are called for. Such changes invalidate, in both contexts, many of the most interesting Bayesian theorems, such as the celebrated "washing out the priors" result. The important point that we wish to stress is that once MB becomes unrealistic (or outright infeasible), AA (together with other potential types of probabilistically incoherent behviour) becomes plausible.

It is worth emphasizing that standard Bayesianism, in either MB or CB form, does not allow for 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 Baysianism: 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 language $\mathcal{L}$. For simple languages, it is reasonable to assume that the agent can comprehend the few logical entailments that these languages contain. And this explains why we entertain the view that MB is a realistic strategy 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 languages), in which a (normatively preferable) MB strategy is attainable.

The paper is organized as follows: Next section describes Schmeidler's two-coin thought experiment that is supposed to produce AA. It also defines formally two of the central concepts of the present paper, namely MB and CB. Section III is an analytic discussion of the arguments for MB mentioned above. It also elaborates on the connection between MB and small worlds, arguing along the lines presented above, for the efficiency of counterfactual probabilistic reasoning in the case that the number of probabilistic hypotheses under consideration is small. Section IV, demonstrates formally the absence of AA under MB reasoning within Schmeidler's two-coin example. Section V concludes.

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

Schmeidler (1989) illustrates the points mentioned above by means of the following coin example, which aims at conveying the same message with Ellsberg's (1961) two-urn paradox: Assume that the agent S considers two coins A and B. S 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 S 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 looses 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 offers her 1$ if the event

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

materializes, and -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" that (given she is less willing to bet on D_H^B than D_H^A and on D_T^B than D_T^A) that

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

and

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

Assume that $P(D_H^B) = P(D_T^B) = 0.4$. It is easy to show that such a P is non-additive. Indeed,

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

and

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

where S is the relevant sample space, namely

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

Hence,

$$1 = P\left(D_H^B \cup D_T^B\right) \neq P(D_H^B) + P(D_T^B) = 0.8 \tag{1}$$

What causes the violation of the additivity property in P? As already mentioned, it is the presence of asymetric information, namely specific information (I_S) about the chance properties of coin A and the lack of similar information for coin B that is responsible for the no-standard properties of P. To be more specific, it is not merely the presence of I_S that caused the non-additivity of I_S, but rather the fact that the agent allowed I_S to affect directly her probabilistic beliefs, instead of utilizing 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. How can such a probability function may be obtained? Simply by the MB counterfactual strategy outlined above, namely by assuming that the agent did not know I_S and contempalted what her probabilistic beliefs would have been in such a counterfactual epistemic state.

After a rather long discussion involving the concepts of MB and CB, it is now time to define these concepts 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 (or $t+\delta$ to be more precise). A CB agent carries out this task as follows:

Classical Bayesianism (CB):

$$P_{t,I_B,I_S,E}(A) = P_t^{I_B,I_S}(A \mid E) \tag{2}$$

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 the point, t_0, in 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):

$$P_{t,I_B,I_S,E}(A) = P_0^{I_B}(A \mid I_S, E) \tag{3}$$

Easwaran (200?) 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." (200?, pp.147). In a similar fashion, Meacham (2003) 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.)" (2003? sleeping beauty paper - emphasis added). Along the same lines, Titelbaum and Kopec (200?), refer to the concept of hypothetical prior as follows: "We can think of an agent's hypothetical prior as representing her "evidential standards" - antecedent to the influence of any contingent evidence, the hypothetical prior encodes how an agent would respond to any package of evidence she might encounter, and which bodies of evidence she would take to support which hypotheses." (200?, pp. 5 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 the 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 (3) in forming her current beliefs, while the alternative actual (and intuitively appealing) course (2) 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)

Before we present our arguments in favor of MB over CB, it is worth emphasizing that despite the term "modern", MB is not a form of Bayesianism that was put forward only recently. Instead, several authors from the 1970s entertained the idea of "hypothetical priors", that is "priors without evidence". For 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).

After this brief (and by no means exhaustive) survey we can now proceed in analyzing the arguments in favor of Modern Bayesianism.

3.1 Permanent Dispositions versus Momentary Inclinations as Determinants of Rational Degrees of Belief.

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, $E_2, ..., E_n$ 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

$$K_n = \cap_{i=1}^{n} E_i. \tag{4}$$

Then Carnap contemplates X's epistemic state at time T_0 in which X posseses 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 Cr_n for a time point T_n is uniquely determined by Cr_0 and K_n : For any H,

$$Cr_n(H) = Cr'_0(H \mid K_n), \tag{5}$$

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, $E_2, ..., E_n$, the proposition K_n, and a function Cr_0 such that (a) E_1, $E_2, ..., E_n$ are possible observation data, (b) K_n is defined by (4), (c) Cr_0 satisfies all requirements of rationality for initial credence functions, (d) the application of (5) 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, $E_2, ..., E_n$, 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 efect of the data E_1, $E_2, ..., E_n$." (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 $Cr_0(\cdot \mid E)$. 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).

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

One of the arguments for MB is akin to a fundamental relation between objective probability (chance) and subjective probability (credence) known as the Principal Principle (PP). PP, originally suggested by David Lewis (1970?), aims at capturing the following intuitive idea: If the agent S knows (with certainty) that the objective probability of the event (proposition) A, $Ch_t(A)$, is x and does not possess any inadmissible evidence for A (more on that below), then S's subjective probability of A must be set equal to x. The question which naturally arises is how to describe formally the aforementioned relationship. One option is to use the current credence function $P_t^{I_B,I_S}$. In such a case, PP takes the form:

$$P_t^{I_B,I_S}(A \mid < Ch_t(A) = x >, E) = x \text{ if } E \text{ is admissible at } t \tag{6}$$

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." (Meecham 2007, pp. 18).

The first version of PP, given by (6), may be interpreted in the light of CB, as follows: The agent S, 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

$$P_{t,I_B,I_S,E}(A) = P_t^{I_B,I_S}(A \mid E).$$

However, if in addition to E the agent learned that the chance of A is x, then

the latter piece of information would dominate the formation of her beliefs, in the sense that $< Ch_t(A) = x >$ screens-off any admissible evidence, E, for A, thus yielding (6). Hence, her new (updated) probability of (A) is given by,

$$P_{t,I_B,I_S,E,<Ch_t(A)=x>}(A) = P_t^{I_B,I_S}(A \mid < Ch_t(A) = x >, E) = x. \qquad (7)$$

A second option for the agent is to operate as an MB, thus treating I_S not as actual but rather as countergactual. 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:

$$P_0^{I_B}(A \mid < Ch_t(A) = x >, E, I_s) = x \text{ if } E \text{ (and } I_s) \text{ is admissible at } t. \qquad (8)$$

What is the restriction that (8) imposes on S's current credences? Put differently, assume S is currently at period t, in which she has come to know I_S, $< Ch_t(A) = x >$ and E. The question is what her current credence $P_{t,I_B,I_S,E,<Ch_t(A)=x>}$ should be. MB's condition (3) implies the following relation:

$$P_{t,I_B,I_S,E,<Ch_t(A)=x>}(A) = P_0^{I_B}(A \mid < Ch_t(A) = x >, E, I_s) = x. \qquad (9)$$

By comparing (7) with (9), 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, Meecham (2007) argues that this is not the case. Specifically, he argues that (6) as oposed to (8) leads to contradictions. To understand why, we must first emphasize the difference in the role that I_S plays in (6) versus (8): In the former case, 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, it is not covered by the admissibility clause - only conditioning information may (or may not) be admissible. Now, let us consider the special case in which A is included in I_S. Since admissibility is not relevant, we cannot claim that A is inadmissible. Hence, (6) dictates that the agent's subjective probability of A is $x < 1$. 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:

$$1 \neq x = P_t^{I_B,I_S}(A \mid < Ch_t(A) = x >, E) = 1$$

On the other hand, the MB formulation (8) of PP does not suffer from this problem. Indeed, since I_S is conditioning information, the admissibility clause applies. Hence, given that we have assumed that A is included in I_S, the latter becomes automatically inadmissible, In such a case, PP loses its force which in turn implies that there is nothing to cause the agent's probability of A to be set equal to $x < 1$.

The preceding analysis has shown that if the agent wishes to secure herself against the possibility to run into contradictions, she must subscribe to MB rather than CB. In a more abstract sense, the purpose of PP is to impose constraints on rational belief in addition to coherence. These constraints aim at

bringing subjective beliefs in line with objective features of the world, namely objective probabilities (in the case that they are known). Are there any other theories of rational beliefs that impose constraints on subjective probabilities over and above coherence. The answer is affirmative. Carnap's system of inductive logic analyzed in the previous subsection is one such case. Of course, Carnap's constraints are quite different than those implied by PP. Specifically, these constraints are introduced in the form of the so-called axioms of symmetry, regularity and convergence. Regardless of their differences, however, both types of constraints should be applied to the prior rather than current probability function, because only for the prior, as Carnap puts it "can we find a sufficient number of requirements of rationality" (1962, pp. 312). 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.

We now turn our attention to the third argument that we propose in favor of MB, namely how to treat the problem of "old evidence".

3.3 Modern Bayesianism and the Problem of Old Evidence

The "problem of old evidence" concerns the role of evidence, e, that was already known at the time of the formation of a scientific theory, T, (hence, it is old) for the confirmation or corroboration of T. This problem highlights a conflict between the (standard) Bayesian theory of confirmation (according to which e does not increase our confidence in T) and the standard scientific practice which assigns to e an important confirmatory role. 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 millenia.... 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 that 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 is known before theory T is introduced at time t. Because e is known at t, $Prob_t(e) = 1$. Further, because $Prob_t(e) = 1$, the likelihood of e given T, $Prob_t(e, T)$, is also 1. We then have:

$$Prob_t(T, e) = \frac{Prob_t(T)Prob_t(e, T)}{Prob_t(e)} = Prob_t(T)$$

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).

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 folllows: 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*

$$T \vdash e \tag{10}$$

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 (10) itself rather than conditionalization on just e. The proposition (10) 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 (10) at t results in an increase in probability of T. Note that this line of defence is based on the assumption that the agent conditionalizes on the entailment (10), 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 (10) (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 defence against the problem of old evidence leads to another problem (potentially more serious) which is incoherence of the agent's probability function steming from logical non-omniscience (more on this below). As will be discussed below, Garber's (1983) solution to the problem of old evidence is based on the idea of weakening the axioms that P_0 has to obey in order to allow for logically non-omniscient rational agents (see, also Hacking 1967 for similar ideas outside the "old evidence" framework).

The second argument for solving the problem of old evidence is more in the spirit of MB since it is based on a 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. 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).

In our set up, the solution outlined above 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".

3.4 Modern Bayesianism and Psychological Detachment from Actual Evidence

Let us now turn our attention to the fourth reason for prefering 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

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

The argument presented above suggests that the subjective evaluation of probabilities should be made, as Longino (1979, pp. 35) put it, "in rational, cool, moments". As already mentioned, such a "cool" and "imparial" agent is referred to as MB in the present paper. As will be shown in the next section, it is the opposite case, namely the one in which the specific information is not treated as conditioning information but rather as a direct determinant (on a par with background information) of the prior probability function (referred to as CB in the present paper) that generates the AA paradox.

3.5 Objections: Logical Non-Omniscience. Local vs. Global Bayesianism and Small vs. Large Worlds

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 $P_0(A \mid E)$ 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 language 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.

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 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 L, capable of expressing all possible hypotheses, all possible evidence, capable of doing logic, mathematics*, etc. In short, 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 phycist 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 language, $\mathcal{L}$ at t_0 was simply not rich enough to accomodate 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. This generates the following question: 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).

More formally, assume that the investigator is interested in the set of specific hypotheses (theoretical propositions) H_i, $i = 1, 2, ..., n$ that cover full set of possibilites 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, $j = 1, 2, ..., m$ 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 relative language" $\mathcal{L}$ 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 language $\mathcal{L}$ rather than the "global" ideal language $\mathcal{L}^*$. 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 Bayseniasm more realistic than global Bayesianism (Bayesianism with a human face, as Jeffrey (1983) put it).

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 relations) 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 entailment relations) that are relevant for the problem in 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*, is it 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 that 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 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 extention as "ridiculous" and "preposterous". Other personalists have suggested a similar course of action. As Teller (1975) observes "...personalists recommend that one simply ignore states, actions, and consequences, the possibility of which have not occurred to the agent or *which seem, prima facie, to be irrelevant to the problem.*" (1975, pp. 170, emphasis added).

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 language, ($\mathcal{L}$ or $\mathcal{L}^*$). Whether MB or CB is appropriate depends on which epistemic status, GB or LB characterises 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 the previous sub-sections.

(iv) CB-LB: Inappropriate or Inefficient Solution.

How does the above classification bear upon AA? The preceding discussion has made clear that the presence of information asymmetry between $\mathcal{F}_k$ and $\mathcal{F}'_k$ is only a necessary condition for AA. A further condition is that the problem in hand falls into the category of large worlds. Put differently, asymmetric information in small worlds is not capable of producing AA. Let us clarify this argument, which forms our basic thesis in this paper. Asume that the agent experiences an information asymmetry at t of the type discussed above, that is at t she 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 informational-symmetrical way. To this end, the simplicity of the problem makes the implementation of the "principle of indifference" or its successor, namely the "maximum entropy principle" perfectly possible. 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 commitements 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 does not work) we do argue that AA has been poorly motivated.

4 Small Worlds, Counterfactual Strategy and Ambiguity Aversion

Let us now analyze Schmeidler's two-coin example discussed in the introduction in the light of the analysis presented in the previous section. To do this, we must first describe the agent's epistemic background at t_0. For simplicity we assume that concerning the coin A, the agent knows with certainty that one of only 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. Put differently, the agent knows at t_0 that the objective probability distribution for coin A is either P_1^A or P_2^A or P_3^A. Similar assumptions are made for coin B, namely the agent knows that one of the following three hypotheses is true,

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

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

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

which give rise to the objective probability distributions P_1^B, P_2^B, and P_3^B, respectively (with $P_2^B(H) = P_3^B(T) = 0.6$).

Now, the agent has to decide about her prior probabilities of the hypotheses mentioned above. Since there is no direct information supporting one hypothesis over the rest, the agent is likely to subscribe 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 that causes an 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 in the introduction. Assuming a minimal degree of probabilistic sophistication on the part of the agent, we have,

$$\begin{aligned} P_0(D_H^A) &= P_0(D_H^A \mid \mathcal{H}_1^A)P_0(\mathcal{H}_1^A) + P_0(D_H^A \mid \mathcal{H}_2^A)P_0(\mathcal{H}_2^A) + P_0(D_H^A \mid \mathcal{H}_3^A)P_0(\mathcal{H}_3^A) = \\ &= 0.5 \times \frac{1}{3} + 0.6 \times \frac{1}{3} + 0.4 \times \frac{1}{3} = 0.5. \end{aligned}$$

In a similar fashion, we obtain,

$$P_0(D_T^A) = 0.5$$

$$P_0(D_H^B) = 0.5$$

and

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

So far so good, Now assume that at time t, $t > t_0$ the agent acquires an important piece of specific information for the problem in 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 the 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 that

$$\begin{aligned} P_t(D_H^A) &= P_0(D_H^A \mid E_A^\infty) = \\ &= P_0(D_H^A \mid \mathcal{H}_1^A, E_A^\infty)P_0(\mathcal{H}_1^A \mid E_A^\infty) + P_0(D_H^A \mid \mathcal{H}_2^A, E_A^\infty)P_0(\mathcal{H}_2^A \mid E_A^\infty) \\ +P_0(D_H^A &\mid \mathcal{H}_3^A, E_A^\infty)P_0(\mathcal{H}_3^A \mid E_A^\infty). \end{aligned} \tag{11}$$

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

$$\begin{aligned} P_0(D_H^A &\mid \mathcal{H}_1^A, E_A^\infty) = P_0(D_H^A \mid \mathcal{H}_1^A) \\ P_0(D_H^A &\mid \mathcal{H}_2^A, E_A^\infty) = P_0(D_H^A \mid \mathcal{H}_2^A) \\ P_0(D_H^A &\mid \mathcal{H}_3^A, E_A^\infty) = P_0(D_H^A \mid \mathcal{H}_3^A). \end{aligned}$$

Furthermore, given that the evidence E_A^∞ is assumed to be arbitrarily large,

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

The above relationships together with (11) imply that

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

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.

$$\begin{aligned} P_t(D_H^B) &= P_0(D_H^B \mid E_A^\infty) = \\ &= P_0(D_H^B \mid \mathcal{H}_1^B, E_A^\infty)P_0(\mathcal{H}_1^B \mid E_A^\infty) + P_0(D_H^B \mid \mathcal{H}_2^B, E_A^\infty)P_0(\mathcal{H}_2^B \mid E_A^\infty) \\ +P_0(D_H^B &\mid \mathcal{H}_3^B, E_A^\infty)P_0(\mathcal{H}_3^B \mid E_A^\infty). \end{aligned} \tag{12}$$

As before, PP implies

$$\begin{aligned} P_0(D_H^B &\mid \mathcal{H}_1^B, E_A^\infty) = P_0(D_H^B \mid \mathcal{H}_1^B) \\ P_0(D_H^B &\mid \mathcal{H}_2^B, E_A^\infty) = P_0(D_H^B \mid \mathcal{H}_2^B) \\ P_0(D_H^B &\mid \mathcal{H}_3^B, E_A^\infty) = P_0(D_H^B \mid \mathcal{H}_3^B). \end{aligned}$$

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 $\mathbf{H}_A = \{\mathcal{H}_1^A, \mathcal{H}_2^A, \mathcal{H}_3^A\}$ and $\mathbf{H}_B = \{\mathcal{H}_1^B, \mathcal{H}_2^B, \mathcal{H}_3^B\}$ the answer is negative. Hence,

$$\begin{aligned} P_0(\mathcal{H}_1^B \quad | \quad E_A^\infty) &= P_0(\mathcal{H}_1^B) = \frac{1}{3} \\ P_0(\mathcal{H}_2^B \quad | \quad E_A^\infty) &= P_0(\mathcal{H}_2^B) = \frac{1}{3} \\ P_0(\mathcal{H}_3^B \quad | \quad E_A^\infty) &= P_0(\mathcal{H}_3^B) = \frac{1}{3}. \end{aligned}$$

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 functions. More specifically, no violations of the additivity property, such as those discussed in the introduction, are observed.

Why did the preceding analysis not produce AA? Put differently, why did the acquisition of E_A^∞ not make the agent to end up with 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 time 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 "went back in time" at point t_0 examined the probabilistic commitements, P_0, that had made at that point (in an epistemic state of not knowing with certainty that E_A^∞ has occured) and stuck to those commitements at t. In this setting, the only way for the agent to exhibit AA is to abandon her commitements made at t_0 and assign 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 from the scene so that there won't be any basis for comparing the probabilistic comitements at t_0 with those adopted at t. 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). 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, the two-coin case analyzed above is quite different from the wildcatter's problem. Consider the following question: 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. The information that the objective probability distribution of coin A is 50-50 is not surprising information at all. In fact the agent had anticipated such a case as indicated by the inclusion of $\mathcal{H}_1^A$ in agent's epistemic language $\mathcal{L}$ at t_0. As a result, the MB strategy is perfectly feasible; in fact it follows quite naturally.

5 Conclusions

(In Bullet Points)

1) There are two ways to view specific information, I_S, as a determinant of the agent's subjective distribution, P_t, at time t: (i) As conditioning information. This pre-supposes the existence of a subjective distribution, P_0, prior to receiving I_S, which plays the role of the vehicle through which I_S affects P_t, via $P_t(\cdot) = P_0(\cdot \mid I_S)$. (ii) As background information on a par with any other background information I_B. In this case, there is no P_0 prior to t, and P_t plays the role of the prior distribution for determining the posterior distributions for times after t, that arise through the acquisition of future information I_f, via $P_{t+1}(\cdot) = P_t(\cdot \mid I_f)$.

2) Are there any reasons, theoretical or practical that favor one interpretation over the other? The answer is yes. There are good reasons, analyzed in detail in Section II, for prefering the first interpretation (MB) over the second (CB).

3) Does the choice of MB versus CB have implications for the main issue under scrutiny in the present paper, namely AA? The answer is yes. If we choose MB there is no room for AA. Being ambiguity averse in such a framework is equivalent to being dynamically inconsistent. On the other hand, under CB, AA can arise because only when the prior distribution is formed at the same time at which I_S arrives informational asymmetries of the type discussed in the introduction can emerge.

4) 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 considerably richer than the one at the time prior to the arrival of I_S.

5) 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 ones is clearly poor since these examples

are textbook cases of small worlds.