



**DEPARTMENT OF INTERNATIONAL AND
EUROPEAN ECONOMIC STUDIES**

ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

**COMPARATIVE IGNORANCE AS AN
EXPLANATION OF AMBIGUITY AVERSION
AND ELLSBERG CHOICES: A SURVEY WITH A
NEW PROPOSAL FOR BAYESIAN TRAINING**

PHOEBE KOUNDOURI

NIKITAS PITTIS

PANAGIOTIS SAMARTZIS

Working Paper Series

24-08

March 2024

Comparative Ignorance as an Explanation of Ambiguity Aversion and Ellsberg Choices: A Survey with a New Proposal for Bayesian Training

Phoebe Koundouri*[†] Nikitas Pittis[‡] Panagiotis Samartzis*[§]

Abstract

Ellsberg-type choices (Ellsberg's paradox) are evidence against the Bayesian theory of Subjective Expected Utility Maximization (SEUM). These choices reflect a particular attitude of the decision maker (DM), namely Ambiguity Aversion (AA). There are two competing interpretations of AA. The first recognizes AA as rational behavior, while the second views AA as a manifestation of a psychological fallacy. This paper focuses on the second interpretation of AA and specifically discusses the most important psychological explanation of AA that has been proposed in the literature, namely Fox and Tversky's (1995) Comparative Ignorance Hypothesis (CIH). CIH holds that AA is mainly a "comparative effect" that occurs when DM feels that he is epistemically inferior for some events of interest compared to others (for which she believes to be epistemically superior). As a result, DM exhibits an aversion towards betting on the epistemically inferior events. The purpose of the paper is twofold: First, to provide a survey of the literature on CIH. Second, to propose a novel "Bayesian Training" (BT) procedure based on "counterfactual thinking". A decision maker who finds BT attractive is likely to move out of the state of comparative ignorance, thereby ceasing to exhibit AA and joining the Bayesian camp.¹

Keywords: counterfactual priors, ambiguity, ellsberg paradox.

JEL Classification: C44, D81, D83, D89

*These authors acknowledge funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 951424; ERC-SyG Water-Futures).

[†]School of Economics and ReSEES Laboratory, Athens University of Economics and Business; Sustainable Development Unit, Athena Research Center; World Academy of Arts and Science; UN SDSN-Europe; e-mail: pkoundouri@aueb.gr, phoebe.koundouri@icre8.eu

[‡]Department of Banking and Financial Management, University of Piraeus, Greece, e-mail: npittis@unipi.gr

[§]Department of Economics, University of Macedonia, psamartzis@uom.edu.gr (corresponding author)

1 Introduction

It is widely accepted that the so-called Ellsberg's paradox poses a serious threat to the empirical validity of the "Subjective Expected Utility Maximization" (SEUM) theory (see Ellsberg 1961). Specifically, Ellsberg-type choices violate one of Savage's (1954) axioms for rational preferences, namely his illustrious Sure Thing Principle (STP)². As a consequence, the decision maker (DM) who exhibits Ellsberg-type choices is not probabilistically sophisticated and does not adhere to SEUM criterion of choice. Lack of probabilistic sophistication means that the degrees of DM's probabilistic beliefs (subjective probabilities or credences) are not *coherent*, that is they are not represented by a unique subjective probability function that obeys the axioms of Kolmogorov (one of which is the axiom of additivity - finite or countable). Put differently, such a DM is unable to come up with a unique numerical probability for each and every event/proposition concerning the decision problem under consideration.

What is the "degree of rationality" of such a probabilistically non-sophisticated / non SEU-maximizer DM? Should be DM condemned as "non-rational" when she exhibits Ellsberg-type behavior? Or alternatively, is the requirement of probabilistic sophistication too strict to be treated as part of the relevant normative ideal? In other words, is it possible that DM's inability to express probabilistically sophisticated beliefs is (in some cases) justified, and as such, it should not be interpreted as evidence against DM's rationality?

One interesting case, pertaining to the aforementioned distinction is that of "Ambiguity Aversion" (AA). A particular case of AA is (roughly) when DM prefers to bet on events/propositions with known probabilities than on those with unknown ones. The question now becomes: Does AA reflect a pathological choice of a non-rational DM, or should AA be included in DM's normative system of preferences? In the latter case (as opposed to the former one), AA may be thought of as an acceptable property of DM's preferences, on a par with other such properties, such as completeness and transitivity. Many authors adopt this more tolerant definition of rationality and proceed in developing a system of preferences in which AA plays a prominent role. Authors following this route argue that AA is not a "behavioral fallacy" arising from DM's misunderstanding of the features of the decision problem at hand. Instead, AA should be seen as part of our pre-theoretical intuitions of "good" decision making. More specifically, apart from the utilities of outcomes and the subjective probabilities of the corresponding events, a third epistemic feature, called "ambiguity", *should* enter DM's decision making process. Hence, the formal representation of these intuitions, in the form of a axiomatic system of preferences, should somehow accommodate AA. Ellsberg writes: "Yet the choices themselves do not appear to be careless or random. They are persistent, reportedly deliberate, and they seem to predominate empirically; many of the people who take them are eminently reasonable, and they insist that they want to behave this

²The "sure-thing" principle, also known as the "independence" principle states that if two actions, f and g entail the same outcome x under the state of nature s , then the ordering of f and g should be "independent" of the value of x .

way, even though they may be generally respectful of the Savage axioms." (1961, p. 656). According to Ellsberg, people do not violate Savage's postulates "by mistake," or "out of a lack of sufficient understanding" but instead violate them "intentionally" because they think it is the rational thing to do.

Axiomatic systems that account for AA have been proposed (among many others) by Schmeidler (1989) and Gilboa and Schmeidler (1989) (see Gilboa and Marinacci 2013 and Machina and Siniscalchi 2014 for surveys of this literature). A common element of these systems is the abolition of STP and the addition of alternative axioms which state explicitly what it means for DM's preferences to display ambiguity aversion. Each of these systems explains Ellsberg's paradox in the sense that within the system, Ellsberg's choices become predictable. Moreover, each system spells out the type of (non-sophisticated) probabilistic beliefs that an ambiguity averse DM possesses. In Schmeidler (1989) and Gilboa and Schmeidler (1989), for example, these beliefs are represented by a convex (non-additive) probability function and by multiple probability functions, respectively. It is important to note that although such probabilistic beliefs are technically referred to as non-sophisticated, they are not irrational.

Contrary to the aforementioned interpretation, another strand of the literature (consisting mainly of strict Bayesians) views AA as "bad pre-theoretical intuition" and as such should not be encoded in the axiomatic formulation of preferences. According to this view, the axioms of rational choice are on a par with the rules of logic or the axioms of arithmetics, meaning that anyone who understands the axioms will feel compelled to obey them. Hence, any person exhibiting AA is arguably non-rational, and thus needs Bayesian training (see, for example, Raiffa 1961, Roberts, 1963, Savage 1972, and more recently Al-Najjar and Weinstein 2009). Moreover, in addition to the aforementioned theoretical reasons, there are also pragmatic reasons why AA is inconsistent with rationality. Indeed, AA together with the resulting incoherence of DM's subjective probability function make her vulnerable to a Dutch book. The latter is a series of bets (offered to DM by a cunning bettor, not better informed than DM) that DM is willing to accept individually, but which jointly inflict upon her a sure loss. A rational DM, so the argument goes, would never accept a set of bets that deterministically lead to her losing money. Furthermore, Al-Najjar and Weinstein (2009) argue that allowing ambiguity attitudes to be "a matter of taste" may account for Ellsberg's paradox but at a prohibitively high cost: AA, viewed as part of rationality standard, "fundamentally contorts the concepts of beliefs and updating, and ends up creating more paradoxes and inconsistencies than it resolves." (2009, p. 250). On this view, if DM exhibits AA then she is not rational and needs probabilistic education. Raiffa (1961) emphasizes the need for teaching people how to act according to Savage's postulates as follows: "If most people behaved in a manner roughly consistent with Savage's theory then the theory would gain stature as a descriptive theory but would lose a good deal of its normative importance. We do not have to teach people what comes naturally. But as it is, we need to do a lot of teaching." (1961, p. 691). In a similar tone, Curley et. al. (1986) refer to AA as a psychological mistake: "According to this 'mistake hypothesis,' in expressing a preference for the less

ambiguous option, subjects are making *a nonconscious, systematic error which, if sufficiently understood, they would correct.*" (1986, p. 233, emphasis added). Savage (1962) characterizes AA as "an unpleasant truth" which, nonetheless, is not "curable by a new theory."

Assuming that Savage's theory remains the normative ideal of decision making, the following question arises: What causes people to exhibit such "deviant behavior"? In other words, what is the root of AA?³ The first studies analyzing the psychological roots of AA appeared in the literature almost the same time with Ellsberg's paper, that is in the beginning of 1960s (see, for example, Roberts 1963, Becker and Brownson 1964, Toda and Shuford 1965, Einhorn and Hogarth 1985, 1986, Curley et. al. 1986). However, the study that attracted the greatest theoretical and empirical interest was that of Fox and Tversky (1995). This study put forward the "Comparative Ignorance Hypothesis" (CIH), as an explanation of AA. The main tenet of this hypothesis is that AA is caused by DM's preference for actions about which she feels she has more knowledge than for those about which she feels ignorant: "Thus, ambiguity aversion represents a reluctance to act on inferior knowledge, and this inferiority is brought to mind *only through a comparison* with superior knowledge about other domains or of other people." (Fox and Tversky 1995. p. 599, emphasis added). More specifically, suppose that DM is interested in the events/propositions of the space \mathcal{F} . Suppose also that \mathcal{F} is partitioned in the the following two sub-spaces \mathcal{F}_1 and \mathcal{F}'_1 , for which DM feels (for some reason) that she has "superior" and "inferior" knowledge, respectively. It is important to emphasize that CIH insists that DM's feeling of inferior knowledge about \mathcal{F}'_1 arises in her mind only in the presence of \mathcal{F}_1 . In other words, if the "superior-knowledge" sub-space \mathcal{F}_1 were absent, and DM encountered only \mathcal{F}'_1 , she would not feel ignorant about it. In the context of CIH, ignorance means "comparative ignorance". Fox and Tversky are quite explicit on this point: "Moreover, we argue that *this contrast between states of knowledge is the predominant source of ambiguity aversion.* When evaluating an uncertain event *in isolation*, people attempt to assess its likelihood - as a good Bayesian would - paying relatively little attention to second-order characteristics such as vagueness or weight of evidence. However, when people *compare* two events about which they have different levels of knowledge, the contrast makes the less familiar bet less attractive or the more familiar bet more attractive." (1995, p. 587-588, emphasis added).

What generates the partition $\{\mathcal{F}_1, \mathcal{F}'_1\}$ in DM's mind? In other words, what are the reasons for which DM might feel that she knows more about \mathcal{F}_1 than about \mathcal{F}'_1 ? To answer this question, we must distinguish between the following two decision making settings (see Heath and Tversky 1991): (i) Chance setups and (ii) Real-world situations. Chance setups, such as "Ellsberg's urns" described in the next section, leave no room for DM to feel "knowledgable" in the specific subject matter, thus mitigating context-sensitivity. There is no meaningful sense in which one can think of herself as an expert on "Heads" or

³If AA is viewed as part of the new non-Bayesian rationality standard, no such question arises.

"Tails" or on "Black" versus "Yellow" balls in an urn. Therefore, the aforementioned distinction is brought about exclusively by the presence of *asymmetric probabilistic information* about the events/propositions of interest. For example, assume that at $t = 1$, DM receives the probabilistic information I_1 that furnishes the objective probabilities of \mathcal{F}_1 but not those of \mathcal{F}'_1 . Furthermore, assume that DM trusts the objective probabilities of \mathcal{F}_1 more than her own subjective probabilities of \mathcal{F}'_1 . In other words, DM recognizes "chance" as a probabilistically more competent agent than herself. DM is therefore in an asymmetric epistemic state in which (because of I_1) she feels she knows more about a useful feature of \mathcal{F}_1 (namely its objective probabilities) than about the same feature of \mathcal{F}'_1 . This asymmetry generates a contrast in DM's mind between her epistemic competence of \mathcal{F}_1 versus that of \mathcal{F}'_1 . As a result, DM experiences comparative ignorance⁴.

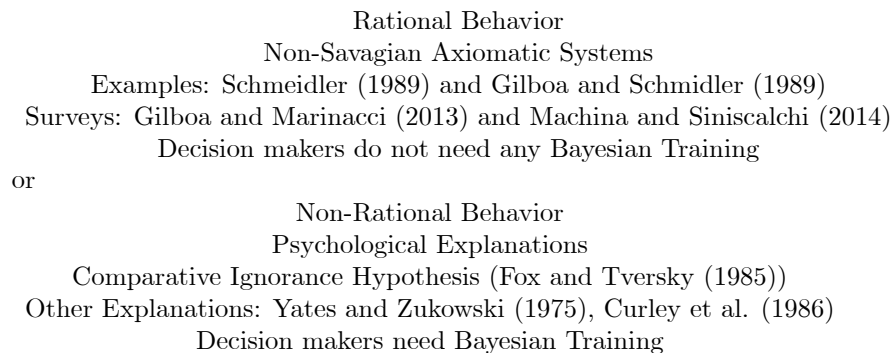
In real-world situations, which are "context-sensitive", the picture is more complex. Asymmetric probabilistic information is not sufficient to produce comparative ignorance. The dominant factor in real-world cases is DM's familiarity or knowledge with the subject matter of each case. If DM feels that she knows more about \mathcal{F}'_1 than \mathcal{F}_1 , then she is likely to bet on the options in \mathcal{F}'_1 than those in \mathcal{F}_1 , despite the fact that \mathcal{F}'_1 is probabilistically vague and \mathcal{F}_1 is probabilistically clear. In these cases, ambiguity refers to the options about which DM feels epistemically incompetent rather than to those which are probabilistically vague.

Next let us assume that DM exhibits AA due to comparative ignorance. The next question is whether DM continues to exhibit the same behavior after having some form of education from a Bayesian trainer (BT). The question is whether BT can convince DM that AA is irrational, or whether DM will ignore BT, thus continuing to see the choice of the non-ambiguous option as the only sensible choice available to her. As already mentioned, many authors in the Bayesian camp believe that AA will cease to exist once DM fully understands the principles of Bayesian rationality. Roberts (1963) advises DM to act as follows: "Put yourself temporarily within the Bayesian framework and see what a careful Bayesian analysis can contribute to your understanding of your answers." (1963, p. 330). How can DM become temporarily Bayesian? The answer to this question takes the form of an educational process, which a Bayesian trainer recommends to DM. Depending on how convincing DM finds the proposed process, she may (or may not) convert to Bayesian orthodoxy. In this paper we propose such a process, implemented in three successive steps, which will be referred to as Bayesian Training Process (BTP). The main goal of BTP is to drive DM out of the epistemic state of comparative ignorance (in which

⁴As mentioned above, an important part of CIH is that if DM encountered only \mathcal{F}'_1 (in isolation) no AA would occur. This point draws a clear distinction between "ambiguity" arising from "comparative ignorance" and "uncertainty" arising from "uniform ignorance". It is rather unfortunate that the two terms, "ambiguity" and "uncertainty" are often used interchangeably in the literature (see Section III for more discussion on this point). This is rather curious since, as will be analyzed in the next Section, in all Ellsberg's examples, from which AA was originated, this comparative feature is explicit.

she is prone to AA), thereby making her form her prior probability function under the state of "uniform" ignorance. A key element of BTP is "counterfactual thinking. In particular, given that the source of comparative ignorance is the presence of the asymmetric probabilistic information I_1 , BTP aims to eliminate comparative ignorance by targeting its source. Specifically, in the first step of BTP, DM is instructed to form her probabilistic beliefs under the (counterfactual) hypothesis that the probabilistic information I_1 is not available and hence, is not certain. In other words, DM's shift to counterfactual thinking aims to temporarily change the modal status of I_1 from "certainty" to "possibility". In this neutral epistemic state, DM is asked to evaluate her subjective probabilities of each proposition A , $A \in \mathcal{F}$ (second step of BTP). Once this task is completed, i.e. DM's prior " I_1 -free" probability function, P_0^c , is elicited, DM is advised to proceed to the third step of BTP, in which she brings I_1 back to the picture by updating her prior beliefs via Bayesian Conditionalization (BC) on I_1 , using P_0^c as the relevant vehicle of conditionalization. If DM accepts each of the above three steps of BTP as reasonable, then it is possible to judge the whole process as reasonable, thus concluding that her previous ambiguity-averse behavior was "wrong".

The foregoing discussion is summarized in the form of the following diagram: Ambiguity Aversion (implies):



This paper focuses on the issues shown on the right side of this diagram. In particular, the rest of the paper is organized as follows: Section 2 analyzes CIH as an explanation of AA. Section 3 surveys the empirical evidence for or against CIH from several psychological studies that have appeared in the literature since the early 80s. Section 4 shows how the proposed counterfactual strategy BTP may drive DM out of comparative ignorance, thus mitigating AA and restoring probabilistic sophistication. Section 5 concludes the paper.

2 The Ellsberg Paradox

Although Ellsberg's paradox is very well known, let us briefly recast it in our own framework and notation and discuss the early evidence of its presence in actual

decision making. The paradox comes in two versions, the "two-colors-two-urns" and the "three-colors-one-urn" ones.

(i) The "two-colors-two-urns" paradox. Consider a DM who is faced with two urns, urn I and urn II. Urn I contains 100 red (R) and black (B) balls in a proportion unknown to the DM. For urn II, the DM is informed that it contains 50 red and 50 black balls. DM is offered the following bets: f_R : "win \$1 if R is drawn from urn I and \$0 otherwise", f_B : "win \$1 if B is drawn from urn I and \$0 otherwise", g_R : "win \$1 if R is drawn from urn II and \$0 otherwise", g_B : "win \$1 if B is drawn from urn II and \$0 otherwise". Ellsberg invites DM to think whether she prefers (i) f_R versus f_B , (ii) g_R versus g_B (iii) f_R versus g_R and (iv) f_B versus g_B . He argues that most decision makers are indifferent between f_R and f_B ($f_R \sim f_B$) as well as between g_R and g_B ($g_R \sim g_B$). However, they tend to prefer g_R to f_R ($g_R \succ f_R$) and g_B to f_B ($g_B \succ f_B$) which implies that, under the fundamental assumption of SEUM theory that DM's beliefs can be derived from DM's betting preferences, DM ends up with non-additive beliefs. To see this, assume that the following equivalences hold:

$$\begin{aligned} g_R \succ f_R &\iff P(R_{II}) > P(R_I) \\ g_B \succ f_B &\iff P(B_{II}) > P(B_I). \end{aligned} \tag{1}$$

Assume that DM's prior subjective probability function, P , is additive with respect to R_{II} and B_{II} , (urn II), that is

$$P(R_{II}) + P(B_{II}) = 1,$$

which implies that

$$P(R_{II}) = 1 - P(B_{II}).$$

Taking into account this equality, the inequalities in (1) become:

$$\begin{aligned} 1 - P(B_{II}) &> P(R_I) \\ P(B_{II}) &> P(B_I). \end{aligned}$$

It immediately follows that

$$1 > P(R_I) + P(B_I)$$

which means that P is not additive with respect to R_I and B_I (urn I). The conclusion is that DM cannot be probabilistically sophisticated with the respect to both urns simultaneously.

(ii) The "three-colors-one-urn" paradox. Consider an urn that contains 90 balls with three different colors. Suppose also, that DM, being at period $t = 1$, is given the specific information I_1 : "30 balls are red and the remaining 60 balls are either black or yellow in unknown proportion". DM will draw a ball at random, which means that each ball has an equal objective probability of being drawn. DM is offered two pairs of choices:

(a) Choose between f and g , where

f : "a bet on red"
 g : "a bet on black".

(b) Choose between f^* and g^* , where:

f^* : "a bet on red or yellow"
 g^* : "a bet on black or yellow".

The following table contains the outcomes for each action and state of nature:

		Outcomes		
		red is drawn	black is drawn	yellow is drawn
$\left[\begin{array}{l} f \\ g \\ f^* \\ g^* \end{array} \right]$	f	100	0	0
	g	0	100	0
	f^*	100	0	100
	g^*	0	100	100

Ellsberg argues that a very frequent pattern of responses is the following:

$$\begin{aligned} f &\succ g \\ g^* &\succ f^*. \end{aligned}$$

This pair of choices implies the following probabilistic relations,

$$\begin{aligned} P(R) &> P(B) \\ P(B \vee Y) &> P(R \vee Y) \end{aligned} \tag{2}$$

where " \vee " denotes the disjunction of the propositions B and Y ². From these inequalities it becomes clear that the additivity property of DM's subjective probability is in trouble. Indeed, since $B \vee Y$ is equivalent to the negation $\neg R$ of R , and $R \vee Y$ is equivalent to the negation $\neg B$ of B , we have

$$\begin{aligned} P(R) &> P(B) \\ P(\neg R) &> P(\neg B). \end{aligned}$$

In such a case, DM thinks of "red" as more probable than "black" and at the same time "not-red" as more probable than "not-black". These pair of inequalities implies a contradiction. More specifically, from (2), we have

$$P(R) + P(B \vee Y) > P(B) + P(R \vee Y)$$

Since B and Y are mutually exclusive, probabilistic coherence requires

$$P(B \vee Y) = P(B) + P(Y).$$

²The initial of each color, e.g. B denotes the proposition "the color in the next draw is Black". Equivalently, it denotes the "event" described in the corresponding proposition.

Similarly,

$$P(R \vee Y) = P(R) + P(Y).$$

The last three relations entail the following,

$$P(R) + P(B) + P(Y) > P(B) + P(R) + P(Y),$$

which is a contradiction.

In both versions of the paradox, DM seems to prefer betting on the events with known probabilities over those judged by DM to be equally probable. In other words, DM displays a tendency to avoid the "ambiguous" or "vague" bet. As will be analyzed below, Ellsberg's concept of "ambiguity" referred exclusively to the case where the probabilities of certain events are vague or ambiguous. Moreover, implicit in Ellsberg's cases is always a comparison between these vague probabilities and the probabilities of some other events which are known to DM. The ambiguity about the event "black", in the "three-color-one-urn" case, settles in DM's mind only when she realizes that there is a similar event, namely "red" for which she knows that its objective probability is 1/3. As mentioned in the introduction, Fox and Tversky's (1995) study clarified this relative aspect of ambiguity aversion (see also Fellner 1961).

2.1 Early Empirical Evidence for Ambiguity Aversion

Shortly after Ellsberg's proposal, several studies empirically tested the "ambiguity aversion hypothesis". Becker and Brownson (1964) provided evidence that in chance setups similar to Ellsberg's, people are willing to pay a premium in order to avoid ambiguity. Moreover this premium is proportional to the *degree of ambiguity*, (defined in the next section). In a similar setting, Yates and Zukowski (1976) gave a characterization of ambiguity in terms of "second-order" probabilities. MacCrimmon (1968) extended the scope of AA from chance setups to real-world situations. A sample of experienced business executives was asked to solve decision problems based on Savage's postulates. The findings indicated that a significant proportion of these agents (approximately 40%) violated Savage's sure-thing principle (STP), thus exhibiting AA. However, the same agents were receptive to Bayesian training as almost all of them revised their initial choices, when challenged by a Bayesian interlocutor. Slovic and Tversky (1974) disputed these results, arguing that the subjects in MacCrimmon's experiment were subject to "subtle pressures" to conform to STP during their discussions with the Bayesian contender. When these pressures are absent, the percentage of subjects that insist on their initial STP violations, and are therefore resistant to Bayesian education, is much higher. Curley et. al. (1984) investigated AA in another natural setting, that of clinical treatments. Their results showed that a small percentage of patients (21%) avoided an ambiguous treatment whose probability of success was estimated to be equal to that of an unambiguous, previously accepted treatment. Hogarth and Kunreuther (1985) provide evidence for AA in insurance business, where the premium people are willing to

pay to insure against an "ambiguous event" is significantly higher than that of a "risky event" (with known probability).

As already mentioned in the introduction, the interpretation of these results is controversial. Some authors (mentioned above) interpret STP violations - particularly persistent violations that are not corrected after Bayesian training - as rational manifestations of AA. Some others disagree, arguing that the above behavior is clearly irrational. Therefore, AA requires further investigation especially regarding its psychological origins. In other words, if AA is the result of the DM's "confusion" about what is a good norm of decision making, then it is important to identify the causes of this confusion. This topic is dealt with in the next section, with particular emphasis on the most prominent of these explanations, the comparative ignorance hypothesis.

3 The Comparative Ignorance Hypothesis: Theoretical Insights and Empirical Evidence

Before we analyze the theoretical underpinning of CIH and discuss its empirical success, let us first give a brief summary of the psychological explanations of AA that preceded CIH.

3.1 Precursors to the Comparative Ignorance Hypothesis

One of the earliest explanations of AA is what Curley et. al. (1986) called the "other-evaluation" hypothesis (see, for example, Fellner 1961, Ellsberg 1963, Roberts 1963). When DM makes a decision, she knows that this decision will be evaluated by others. Hence, she exhibits a preference towards the most justifiable choice. For example, in the context of Ellsberg's "three-colors-one-urn" DM chooses "red" over "black" because she believes that this choice (associated with the *known* objective probability of "red") is easier to justify than the choice "black" (related to the *unknown* objective probability of "black"). A variant of this hypothesis is the so-called "self evaluation" hypothesis, according to which DM does not expect to be evaluated by others but rather by herself in the future (see, for example, Toda and Shuford 1965).

Another explanation was proposed by Yates and Zukowski (1975). These authors argue that AA is due to DM's erroneous prejudice that the outcomes associated with the ambiguous choice are generated from a competitive non-random process that favors the undesired outcomes. Curley et. al. (1986) called this explanation "the hostile nature" hypothesis. For example, in the case of Ellsberg's urn, the outcome "black", associated with the choice g , is produced by an unknown process which nevertheless tends to produce more "yellow" than "black" outcomes. This "psychological bias" is caused by DM's realization that she has no control over generating favorable outcomes (see, Langer 1975).

Einhorn and Hogarth (1985, 1986) proposed an explanation of AA based on the following idea: People form their subjective probabilities on the basis of an "anchoring-and-adjustment process". More specifically, the process starts

with the assessment of an initial probability $P(D) = p_D$ of some event, D , that serves as the "anchor". The final subjective probability of D is calculated as $P(D) + k$ where k is the net effect of the adjustment process for ambiguity. This process involves a mental simulation in which DM imagines higher and lower values of p_D . The size of this simulation, in turn, depends on the amount of DM's ambiguity about p_D . In the limiting case, in which p_D is known, no mental simulation occurs. Besides the amount of ambiguity, this model accounts for DM's attitude towards ambiguity as well. In particular, a parameter of this model controls the relative weights by which probabilities higher and lower than the anchor p_D are weighted. Einhorn and Hogarth conduct a series of experiments whose results seem to support their model.

Curley et al. (1986) conducted a comparative test of which hypothesis, among a set of five alternative hypotheses, best explains AA. The five hypotheses they considered are the following: (i) "other-evaluation", (ii) "hostile nature", (iii) "self-evaluation", (iv) "forced choice" (Roberts 1963), and (v) "uncertainty avoidance" (Curley et al. 1986). Of these five hypotheses, only the first received some empirical support. None of the remaining four hypotheses were found to be empirically relevant. The general consensus in the late 1980s was that although AA was empirically well documented, its origins were poorly understood. This gap in the literature was filled by the comparative ignorance hypothesis, to which we now turn.

3.2 "Comparative Ignorance" and "Competence" Hypotheses

CIH is based on the so-called "Competence Hypothesis" (CH), introduced by Heath and Tversky (1991), according to which: "The willingness to bet on an uncertain event depends *not only on the estimated likelihood of that event* and the precision of that estimate; it also depends on one's general knowledge or *understanding of the relevant context.*" (1991, p. 7, emphasis added). This means that DM displays an aversion towards betting on situations for which she feels incompetent: "We propose that - holding judged probability constant - people prefer to bet in a context where they consider themselves knowledgeable or competent than in a context where they feel ignorant or uninformed." (1991, p. 7). Fox and Tversky introduce CIH by asking the following question: "If ambiguity aversion is driven by the feeling of incompetence ... the question arises as to *what conditions produce this state of mind.*" (1995, p. 587, emphasis added). They answer this question as follows: "We propose that people's confidence is undermined when they *contrast* their *limited knowledge* about an event with their *superior knowledge* about another event, or when they compare themselves with more knowledgeable individuals. Moreover, we argue that this *contrast between states of knowledge is the predominant source of ambiguity aversion.*" (1995, p. 587, emphasis added). These authors conduct a series of experiments all of which lend support to the hypothesis that "the Ellsberg phenomenon is an inherently comparative effect" (1985, p. 600). The overall evidence suggests that AA is significantly reduced or even eliminated when the

"comparative feature" is removed from the decision-making context.

Fox and Tversky argue that CIH explains the Ellsberg choices. Specifically, in the context of the "two-urns-two-colors" case, DM chooses urn II over urn I because she feels that she has inferior probabilistic knowledge about urn I than urn II. In other words, she is comparatively more ignorant about urn I than urn II. Similarly, in the context of "three-colors-one-urn" case, she prefers to bet on "red" than on "black" because she feels that she knows more about "red" than about "black", namely the objective probability of "red" and not that of "black". It is important to emphasize that what brought the distinction between superior and inferior knowledge in DM's mind is the "comparative structure" of the Ellsberg problems: DM analyzes the choice of selecting urn I *jointly* with that of urn II. Similarly, she considers the choice of betting on a red ball *jointly* with that of betting on a black ball.

The aforementioned discussion implies the following: AA and the resulting Ellsberg's choices do not arise because DM failed to elicit a unique and precise probability, $P_1(B)$, of B , $B \in \mathcal{F}'_1$. Instead, AA emerges if DM contrasts her subjective probability of B with the epistemically more reliable (according to her own lights) objective probability of A , $A \in \mathcal{F}_1$. In a paper that may be seen as a precursor of Fox and Tversky's paper, Sahlin (1993) offers the following example that emphasizes this point: "Assume that you are offered two lotteries and your task is to choose the one you consider to be most preferable. The first lottery gives you 100 pounds if you draw a white ball from an urn containing 30 white balls and 70 black balls; otherwise you get nothing. The second lottery gives you 100 pounds if there is a transit strike in Verona, Italy, next week; otherwise nothing. For the sake of argument, assume that, after considering it carefully, you believe that the probability that there will be a transit strike in Verona next week is 0.30. Thus, provided pounds and utilities are exchangeable, the (subjective) expected utility of this lottery is 30 pounds. Thus the second lottery obviously has the same expected value as the first one. But although the two gambles have the same expected utility, you will trade the first gamble for the second and this preference conflicts with the recommendations of the Bayesian doctrine." (1993, p. 14). In this example, DM does not fail to come up with a precise probability of "a strike in Verona next week"; she somehow managed to elicit her subjective probability of 0.30. What made her preferring the first lottery to second is the fact that she feels that she is epistemically more competent about the urn than about the socio-economic factors that might cause a strike in Verona."

The foregoing analysis suggests the following causal chain of events that lead to AA:

- 1) At $t = 1$, DM receives the probabilistic information I_1 (from a fully reliable source) which furnishes the objective probabilities of \mathcal{F}_1
- 2) I_1 naturally suggests the partition $\{\mathcal{F}_1, \mathcal{F}'_1\}$ in DM's mind and invites her to epistemically evaluate this partition
- 3) DM considers jointly (compares) how knowledgeable she is about \mathcal{F}_1 relative to \mathcal{F}'_1
- 4) DM decides that she has superior knowledge of \mathcal{F}_1 relative to \mathcal{F}'_1 (she

trust "chance" more than her own probabilistic judgements

5) DM enters the psychological state of comparative ignorance

6) DM exhibits ambiguity aversion and Ellsberg choices.

At this point the following question arises: Are there cases in which DM bypasses the probabilistic information I_1 , thereby preferring to bet on the vague event B , $B \in \mathcal{F}'_1$? (under the assumption that $P_1(B) = P_1(A | I_1)$)? Heath and Tversky (1991) answer this question in the affirmative, especially for real-world judgmental problems. Specifically, according to CH, there may be cases in which DM is more ignorant about the objective probability of B than that of A , but at the same time she feels more knowledgeable (or skillful) about the subject-matter of B than that of A . Hence, the probabilistic-information deficit of B relative to A is over-compensated by a general-knowledge surplus of B relative to A . This means that the fourth step in the aforementioned causal chain is missing. In the context of the Ellsberg "three-colors-one-urn" paradox, DM could ignore I_1 , thus preferring to bet on "black" rather than "red" if she considered herself an "black-ball expert". As a more realistic example, a DM who consider himself as a "soccer expert" would prefer to bet on the outcome of a soccer match based on his own probabilistic judgement over a chance event of equal objective probability. What is important to emphasize is that the dominant factor in establishing comparative ignorance in DM's mind is not so much the simultaneous existence of probabilistically clear and vague options, as the existence of options about which DM feels epistemically inferior *together with* options about which she feels epistemically superior. Put differently "ambiguity" does not refer (solely) to options that are probabilistically vague but rather to options about which DM feels relatively ignorant.

3.3 A Taxonomy of the Cases Implied by CIH

Let us now summarize the foregoing discussion in the context of a simple organizing framework, based on the concept of second-order probabilities (see, for example, Camerer and Weber 1985):

Let us suppose that at $t = 1$, DM receives the probabilistic information I_1 , informing her about the objective probabilities of \mathcal{F}_1 . The objective probabilities of \mathcal{F}'_1 remain unknown. Furthermore, assume that DM fully trusts the source of information I_1 . This latter assumption is translated into $P_1(I_1) = 1$, that is her subjective probability of the truth of I_1 is equal to unity. Based on I_1 and under the additional assumption that DM subscribes to "the principle of direct inference", DM's subjective probability of A is equal to

$$P_1(A) = p_A$$

³For example, in the context of the "three colors-one-urn" case, if A is the proposition "the next draw is a red ball", $p_A = 1/3$. How firmly does DM

³The "principle of direct inference" is a fundamental rationality principle that relates DM's subjective probabilities with the corresponding objective probabilities. According to this principle, when DM knows the objective probability of A , she should adopt it as her subjective probability of A .

believe the proposition " $P_1(A) = p_A$ "? This question brings us to the realm of "probabilities of probabilities", or "second-order probabilities". If DM fully trusts I_1 , in the sense $P_1(I_S) = 1$, then her subjective (second-order) probability, Q_1 , for " $P_1(A) = p_A$ " at $t = 1$ is equal to one. Specifically,

$$Q_1(P_1(A) = p_A) = 1, \quad A \in \mathcal{F}_1. \quad (3)$$

The last equation implies that all the propositions of \mathcal{F}_1 are characterized by the same degree of epistemic reliability, with this degree being equal to unity.

Next, let $B \in \mathcal{F}'_1$ for which no specific information is available. For example, B is the proposition "the next draw is a black ball". What is DM's subjective probability of B ? Assume that in the absence of any information on the objective probability of B , DM thinks that her subjective probability, $P_1(B)$, of B can take on any of the following n values, $p_{B,1}, p_{B,2}, \dots, p_{B,n}$. How "probable" does she find each of these values? This question is answered by her own second-order probabilities, $q_{B,1}, q_{B,2}, \dots, q_{B,n}$, corresponding to $p_{B,1}, p_{B,2}, \dots, p_{B,n}$, respectively, with $\sum_{i=1}^n q_{B,i} = 1$. Specifically,

$$\begin{aligned} Q_1(P_1(B) = p_{B,1}) &= q_{B,1} \\ Q_1(P_1(B) = p_{B,2}) &= q_{B,2} \\ &\cdot \\ &\cdot \\ &\cdot \\ Q_1(P_1(B) = p_{B,n}) &= q_{B,n}. \end{aligned} \quad (4)$$

Q_1 may be interpreted as DM's *epistemic reliability function*, which shows how reliable DM gauges each of her first-order credence functions to be (see, Gardenfors and Sahlin 1982). Sahlin (1993) interprets the first-order probabilities, P_1 , as "ordinary subjective probabilities, i.e. as probabilities *qua* basis of action" and the second-order probabilities, Q_1 , as "epistemic probabilities i.e. as measures of the quality of knowledge." (1993, p. 26). Now, let us make the additional assumption, that DM is willing to adopt as her (first-order) subjective probability of B the weighted average of $p_{B,i}$, $i = 1, 2, \dots, n$ using $q_{B,i}$, $i = 1, 2, \dots, n$ as the corresponding weights:

$$E(P_1(B)) \equiv \bar{P}_1(B) = \sum_{i=1}^n q_{B,i} p_{B,i}. \quad (5)$$

The last relationship shows that DM can, in principle, calculate a single probability for both $A \in \mathcal{F}_1$ and $B \in \mathcal{F}'_1$, namely $P_1(A)$ and $\bar{P}_1(B)$, respectively. It is worth noting that this solution was proposed by Roberts as early as 1963 (see Roberts 1963, p. 329).

Let us now assume that (a) $P_1(A) = \bar{P}_1(B)$ and (b) A and B produce the same monetary outcomes, e.g. if A is true then DM wins \$100 and zero otherwise and so is the case of B . Now we are ready to define the following cases:

(i) The Bayesian Case: If DM is Bayesian, then she will be indifferent between betting on A and B . She trusts her own probabilistic judgements, summarized by $\bar{P}_1(B)$, to the same extent that she trusts the objective probability $P_1(A)$. In this case "uncertainty" is reducible to "risk" as dictated by the standard Bayesian orthodoxy. DM's ambiguity about her subjective probability of B is not an additional factor (beyond probabilities and utilities) affecting her choices. Becker and Brownson (1964) introduced the concept of "degree of ambiguity" encapsulated (in the context of our framework) by the "standard deviation", σ_{P_B} , of the "random variable" $P_1(B)$. For example, imagine two Bayesian decision makers, X and Y, who agree on the subjective mean of $P_1(B)$, i.e.

$$\bar{P}_1^X(B) = \bar{P}_1^Y(B)$$

but disagree on the standard deviation of $P_1(B)$, for example,

$$\sigma_{P_B}^X > \sigma_{P_B}^Y. \quad (6)$$

According to Becker and Brownson's view, the decision maker X exhibits a "higher degree of ambiguity" about $P_1(B)$ than does Y. However, if this "ambiguity feature" does not affect the betting dispositions of either X or Y, so both are indifferent between betting on A or B , then both X and Y remain Bayesian, regardless of (6).

(ii) The Epistemic Inferiority Case: Here, we assume that DM considers herself as epistemically inferior to chance. In other words, she considers "chance" as an epistemically more competent "agent" than herself. Hence, she feels that she is more ignorant about B than about A which makes her to prefer betting on A relative to B . In this case, the asymmetric information I_1 translates into an aversion towards betting on B over A . This type of ambiguity, arising mainly in "chance setups" will be hereafter referred to as AA1. It is worth emphasizing that between 1961 and 1995, AA1 was the only conceivable type of ambiguity in the literature. "Ambiguity" referred exclusively to vague objective probabilities, and "ambiguity aversion" meant exclusively the avoidance of such vagueness. After 1995, the year Fox and Tversky's study was published, the concept of ambiguity was expanded to include the third type analyzed below.

AA1 implies that DM ends up with two different subjective probabilities of B : The first is $\bar{P}_1(B)$, as calculated by (5). This probability is obtained from DM's "judgement". The second, referred to as $P'_1(B)$, is derived from her "choice", that is from DM's preference to bet on A ($A \in \mathcal{F}_1$) than on B ($B \in \mathcal{F}'_1$). In fact, this choice implies that $P'_1(B) < P_1(A) = \bar{P}_1(B)$, which means that DM forms two probabilities of B , one from her own judgement ($\bar{P}_1(B)$) and one from her own choice ($P'_1(B)$). This constitutes a violation of a fundamental rationality principle of Bayesianism, namely DM's subjective probability function is a) derivable from her preferences and b) unique and coherent. Heath and Tversky(1991) remark: "Moreover, our results call into question the basic idea of defining beliefs in terms of preferences. If willingness to bet on an uncertain event *depends on more than the perceived likelihood of that event* and the confidence in that estimate, it is exceedingly difficult - if not

impossible - to derive underlying beliefs from preferences between bets.." (1991, p. 26, emphasis added). All versions of Ellsberg's paradox can be classified as members of AA1.

In a study partly anticipating CIH, Frisch and Barron (1988) argue that AA1 results mainly from the absence of probabilistic information for \mathcal{F}'_1 : "Ambiguity is uncertainty about probability created by missing information that is relevant and *could be known*." (1988, emphasis added). In fact, what makes DM to think that the missing information could be known is the presence of similar information, namely I_1 for \mathcal{F}_1 that is known. It is worth noting that the comparative context, suggested explicitly by Fox and Tversky (1995), is implicit in Frisch and Barron (1988).

(iii) The Epistemic Superiority Case: This case, hereafter referred to as AA2, refers mainly to "real-world" situations, which are "context-sensitive". Here, the picture is more complex. Asymmetric probabilistic information is not necessary for comparative ignorance. The dominant factor in real-world cases is DM's familiarity or knowledge with the subject matter of each case. If DM feels that she knows more about \mathcal{F}'_1 than \mathcal{F}_1 , then she is likely to bet on the options in \mathcal{F}'_1 than those in \mathcal{F}_1 , despite the fact that \mathcal{F}'_1 is probabilistically vague and \mathcal{F}_1 is probabilistically clear. In these cases, ambiguity refers to the options about which DM feels epistemically incompetent rather than to those which are probabilistically vague. This means that DM prefers to bet on B ($B \in \mathcal{F}'_1$) rather than A ($A \in \mathcal{F}_1$). Again, this epistemic asymmetry results in DM having two subjective probabilities of B , namely $\bar{P}_1(B)$ and $P'_1(B)$, but now with $P'_1(B) > P_1(A) = \bar{P}_1(B)$. DM's impression that she is more ignorant about the chancy events of \mathcal{F}'_1 than about the context-specific events of \mathcal{F}_1 drives her to the state of comparative ignorance, thus producing AA2-type behavior. AA2 may result from what Langer (1975) called, "illusion of control". This psychological state emerges when "...factors from skill situations (competition, choice, familiarity, involvement) introduced into chance situations cause individuals to feel inappropriately confident." (1975, p. 311).

In a series of experiments, Heath and Tversky (1991) documented several instances of AA2. In one of these experiments, participants were asked at first whether they considered themselves experts in two subjects: football and politics. Those who declared that they are football experts were given the following option. To bet on a football match or on the result of a lottery that they themselves considered equally probable. But when the same participants were given the choice to bet on the winner of the presidential elections in 13 US states in November 1998 or on the equiprobable outcome of a lottery, they chose the lottery. Similar results were obtained for those who declared that they are experts in politics. It is interesting to emphasize that what ultimately counts in the manifestation of phenomena of AA2-type is whether DM *feels* that she is epistemically superior on a specific subject and not whether she actually *is*. For example, in the aforementioned experiment, "the strategy of betting on judgment was less successful than the strategy of betting on chance in both data sets. The former strategy yielded hit rates of 64% and 78% for football and election, respectively, whereas the latter strategy yielded hit rates of 73% and

80%." (1991, p. 15).

Fox and Weber (2002) argue that several well-known paradoxical patterns of behavior, such as investors' preference to invest in the domestic stock market rather than in foreign markets (home bias) may be classified as AA2 cases. In all these cases, a source of uncertainty perceived (even erroneously) as familiar by DM (football, politics, domestic stock market) is contrasted with another less familiar source of uncertainty (lotteries, foreign stock market), thus driving DM into the state of comparative ignorance. Fox and Weber believe that CIH "...represents a sharp break from previous accounts of decision under uncertainty because it asserts that decisions are influenced by the cognitive context in which the decision maker finds him or herself so that a particular uncertain prospect may be more or less attractive *depending on whether or not a contrasting state of knowledge is salient.*" (2002, p. 6, emphasis added).

3.4 Ambiguity versus Uncertainty

All the above cases are based on the assumption that DM is willing to take the step from (4) to (5), that is to summarize the second-order probability distribution, Q_1 (for B) with its mean $\bar{P}_1(B)$. Alternatively, DM may be reluctant to make this move, thus preferring to represent her credal state about B by the full distribution Q_1 rather than just its mean. In other words, DM may find it more appropriate to maintain multiple subjective probabilities of B , namely $p_{B,1}, p_{B,2}, \dots, p_{B,n}$ instead of combining them into a single value using (5). In such a case, DM still exhibits non-Bayesian behavior but now its origins are different from those of AA1 and AA2: DM fails to come up with a single, coherent subjective probability function not because of ambiguity aversion, but because her beliefs are inherently "indeterminate". Walley (1991) argues "It seems clear that indeterminacy exists. A little introspection should suffice to convince You that Your beliefs about many matters are presently indeterminate." (1991, p. 210). This means that there are two sources of non-Bayesian incoherent beliefs. One is comparative ignorance which can cause ambiguity aversion and the second is uniform ignorance which can cause indeterminate beliefs. Let us refer to these cases as "ambiguity" and "uncertainty", respectively. As mentioned in the introduction, it is misleading to consider these two cases as equivalent. In our view, Ellsberg's paradox was designed to unearth the (then new) epistemic state, of "ambiguity" (in which the DM knows the probabilities of \mathcal{F}_1 but not of \mathcal{F}'_1). This state lies between the traditional states of "risk" (in which the DM knows the objective probabilities of all the elements of \mathcal{F}) and "uncertainty" (in which the DM does not know any of the objective probabilities of the elements of \mathcal{F}). According to this view, "ambiguity" describes an epistemic state that is distinct from that described by "uncertainty", which means that the two terms should not be used interchangeably. This distinction was clearly articulated as early as 1964 by Becker and Bowson (1964) who defined ambiguity as an epistemic condition "falling between two extremes - 'complete ignorance' and 'risk'" (1964, p. 62). In contrast, Gilboa and Marinacchi (2016) do not find such a distinction meaningful: "Today, the terms "ambiguity", "uncertainty"

(as opposed to “risk”), and “Knightian uncertainty” are used interchangeably to describe the case of unknown probabilities." (2016, footnote 8). But this is not what Ellsberg’s paradox was designed to capture. The problem of how the DM assigns probabilities to \mathcal{F} under uncertainty was well-known long before 1961, year at which Ellsberg devised his paradox. The new situation that Ellsberg’s paradox brought to light is the one in which the DM faces risk and uncertainty *within the same decision problem*, in the sense that she knows the probabilities of \mathcal{F}_1 but not those of \mathcal{F}'_1 . In other words, DM’s simultaneous exposure to risk and uncertainty is the trigger that may cause her to display Ellsberg-type behavior.

In a study that to some extent predicted CIH, Fellner (1961) emphasized the comparative character of AA. He considers two events, E_x and E_y produced by the processes X and Y , respectively. X is a "chance process" with well-defined objective probabilities, such as the tossing of a fair coin. Y , on the other hand, is a "stochastic real-world phenomenon" that produces "chancy events" with vague probabilities, (for example, the stock market). Assume that the objective probability of E_x is identical to DM’s subjective probability of E_y , i.e. $P(E_x) = P(E_y)$. At this point Fellner raises the issue of "comparability" between $P(E_x)$ and $P(E_y)$: "The theory (the standard Bayesian theory) doubtless postulates that probability judgments for various processes *are strictly comparable with each other*. One cuts across processes, so to speak, without fear of distortion." (1961, p. 672 emphasis added). Strict comparability between $P(E_x)$ and $P(E_y)$ means that (under the additional assumption that DM’s utility of the outcome associated with E_x is equal to that associated with E_y) DM is indifferent between betting on E_x and betting on E_y . AA emerges only when after DM has compared the two probabilities, she exhibits a preference towards $P(E_x)$ (despite the fact that $P(E_x) = P(E_y)$). This in turn implies that in the absence of $P(E_x)$, that is when $P(E_y)$ were contemplated in isolation, the issue of "comparison" would never have arisen and the resulting AA would never have occurred.⁴ In fact, the experimental results of Fox and Tversky (1995) confirm exactly this: AA does not arise when the two cases of risk and uncertainty are considered separately. In particular, AA disappears entirely when people make decisions in a non-comparative epistemic context. Fox and Tversky write: "When evaluating an uncertain event in isolation people attempt to assess its likelihood - as a good Bayesian would - paying relatively little attention to second-order characteristics such as vagueness or weight of evidence" (1995, pp. 587-588). The process of "Bayesian training" described in the next section, aims to remove this "comparative feature" from DM’s probabilistic assessments. In other words, it aims to force DM to build her subjective probability function under the epistemic state of "uniform" rather than "comparative" ignorance. Moreover, when DM is in the epistemic state of uniform ignorance, she can invoke a widespread epistemic principle, the Principle of Indifference (POI), in order to construct determinate subjective probabilities (see next section).

⁴Of course, even in the absence of $P(E_x)$, DM might have struggled to come up with a determinate probability of E_y . In this case, she would have once again violated the Bayesian principles, but now for a different reason that has nothing to do with AA.

3.5 More Recent Evidence for CIH

After the publication of Fox and Tversky's study in 1995, several papers appeared in the literature that further tested CIH in various contexts. Chow and Sarin (2001) provide evidence for CIH although not as strong as that of Fox and Tversky. Specifically, they show that when DM moves from the comparative to the non-comparative setting, AA decreases but does not disappear entirely: "The key finding that emerges from our experiments is that the clear bet is priced higher than the vague bet under both comparative and non-comparative conditions. The comparison, however, enhances the difference in prices between clear and vague bets. In the absence of a direct comparison (non-comparative condition) this difference is smaller, but it does not disappear." (2001, p. 138).

Arlo-Costa and Helzner (2005) provide evidence against CIH by showing that AA arises in a non-comparative context as well. For them, AA is a manifestation of peoples' inability to form precise subjective probabilities in the epistemic state of uniform ignorance.

Muthukrishnan et. al (2009) show that AA is the main reason why consumers prefer "established brands", i.e. brands for which consumers believe to be of better quality. Moreover, they provide evidence for CIH by showing that differences in perceived quality between brands are present only in comparative contexts.

Rubaltelli et. al. (2010) attempt to explain CIH in terms of people's "affective reactions". More specifically, they show that people's positive affective reactions to the "clear" bet are stronger (more positive) in a comparative rather than a non-comparative context: "Therefore, the present results strongly suggest that ambiguity avoidance depends on the affective reactions people perceive towards clear/familiar stimuli and ambiguous/unfamiliar stimuli in J(oint) E(valuation)." (2010, p. 253).

Shapiro (2020) invokes the CIH to explain the empirical fact that business training programs have limited or even negative effects on the post-training profits of the participants in these programs. The basic idea is as follows: Participating in such a program exposes the participants to people more knowledgeable than themselves, for example trainers or more successful entrepreneurs. As a result, the participants enter a state of comparative ignorance, thus exhibiting an aversion to the "ambiguity" of the training program.

4 Bayesian Training

Let us now assume that CIH is true, which means that AA is attributable solely to comparative ignorance. As analyzed above, a necessary condition for DM to exhibit comparative ignorance in chance setups, such as those of Ellsberg's urns, is that at the beginning of her epistemic life, namely $t = 1$, she possesses asymmetric information I_1 , which induces a partition $\{\mathcal{F}_1, \mathcal{F}'_1\}$ of DM's algebra of propositions \mathcal{F} . This means that if (counterfactually) at $t = 1$, I_1 were not available, then the aforementioned partition would not have occurred, DM

would not have entered the state of comparative ignorance, thus forming her prior in the state of uniform ignorance.

The aforementioned epistemic situation is different from the standard form of Bayesian Confirmation Theory (BCT). In the context of BCT, DM begins her epistemic life at period $t = 0$, in which she does not possess any probabilistic information, I , about the events of interest \mathcal{F} . This enables her to form her information-free prior P_0 in an epistemic state in which I (any I) is a contingency rather than a reality. Let us hereafter refer to P_0 as "Ur-prior" (see, for example, Meacham 2016). In this state of uniform ignorance, DM evaluates her prior subjective conditional probabilities, $P_0(A | I)$ for every $A \in \mathcal{F}$ and for every possible information scenario I that she may encounter in the future. For example, in the case of Ellsberg's "three-color-urn", DM has to form conditional Ur-priors for the event "red" conditional not only on the information I_1 : "the urn contains 30 red balls", but also on the information I_2 : "the urn contains 40 red balls" as well as on I_3 : "the urn contains 30 red balls or 20 yellow balls", etc. Having completed the task of ascertaining her Ur-prior at $t = 0$, DM is now ready to wait for the probabilistic information to arrive at $t = 1$. Suppose that at $t = 1$, DM receives the information I_1 . Then Bayesian Conditionalization dictates that DM forms her posterior (unconditional) probability P_1 , by following the rule $P_1(A) = P_0(A | I_1)$ for every $A \in \mathcal{F}$. This means that what Bayesian rationality demands of DM is to form P_0 in a manner that is consistent with the rules of mathematical probability only once in her epistemic life, and then "sit back and enjoy the evidential ride" (Strevens 2006, p.9). If DM at $t = 1$ refuses to set $P_1(A)$ equal to $P_0(A | I_1)$, then she is "dynamically inconsistent", and this inconsistency makes her susceptible to a "diachronic Dutch Book" (see, Teller 1973).

The foregoing discussion raises the following question: If the information-free time point $t = 0$ is absent in DM's epistemic problem (as it is in all versions of Ellsberg's paradox), and the beginning of DM's epistemic life coincides with the time she already possesses I_1 , how can DM form an Ur-prior? The answer to this question lies at the heart of the following process of Bayesian training, hereafter referred to as BTP:

Step I: In order to form her I_1 -free prior, DM must first move from the actual epistemic state, \mathcal{E}_1 , in which he (actually) knows that I_1 is true to the counterfactual epistemic state, \mathcal{E}_0 , in which she would be if she were not sure that I_1 is true. To achieve the mental transfer from \mathcal{E}_1 to \mathcal{E}_0 , DM is instructed to temporarily "delete" I_1 from her corpus of knowledge, thereby thinking solely in terms of the existing "background" information (e.g. there is an urn containing 90 red, black and yellow balls of unknown proportions). The important thing to note is that by making this *counterfactual move*, DM brings herself in an epistemic state in which she is equally uninformed about the probabilities of \mathcal{F}_1 and \mathcal{F}'_1 , that is, she restores her "uniform ignorance" over \mathcal{F} . Before proceeding further, the following clarification is in order: When DM is instructed "to delete I_1 ", she is advised to do so only temporarily. No canon of rationality would tell DM to discard useful information. On the contrary, DM should update her beliefs in the lights I_1 by means of BC (see third step).

Step II: Now DM is (counterfactually) in the state of uniform ignorance in which she is asked to form her *counterfactual Ur-prior* P_0^c . In particular, DM is instructed to ask herself the question "what would my prior probability function P_0^c be, if I did not know I_1 ?". By asking herself this question, DM is effectively looking for her Carnapian prior (see Carnap 1962), or as we called it above, her "Ur-Prior". As Meacham (2016) remarks: "This function might be understood in various ways, but common candidates include: the credences (a.k.a subjective probabilities) a subject should have if she had no evidence, a subject's initial credences, a subject's evidential standards, and any function that plays the right diachronic role." (2016, pp. 1). In the epistemic context of \mathcal{E}_0 , all of DM's probabilistic assignments to \mathcal{F} exhibit the same degree of epistemic reliability, and because of this, DM is protected from entering the cognitive state of comparative ignorance. DM can now evaluate all her unconditional subjective probabilities $P_0^c(A)$, $A \in \mathcal{F}$ as well as all her conditional subjective probabilities $P_0^c(A | I_k)$, $A, I_k \in \mathcal{F}$, $k = 1, 2, \dots$. DM is also instructed about how to calculate all these probabilities, in the spirit of BCT. Specifically, she is told that she has to specify a set \mathcal{H} of n theoretical hypotheses, H_1, H_2, \dots, H_n , each of which defines an objective probability function Ch_i , $i = 1, 2, \dots, n$ on \mathcal{F} . Then, DM should assign coherent, prior subjective probabilities $P_0^c(H_1)$, $P_0^c(H_2), \dots, P_0^c(H_n)$ to the hypotheses H_1, H_2, \dots, H_n , respectively. The question is how DM should accomplish this task, or in other words, how she should solve "the problem of the priors". The answer to this question depends on whether BT is a "subjective" or "objective" Bayesian. For the former, any probabilistic assignment to H_1, H_2, \dots, H_n is eligible as long as it is coherent. For the objective Bayesian however, this solution is not admissible. On this alternative view, the prior subjective probabilities $P_0^c(H_i)$ must be objectively determined. The only rational prior distribution, P_0^c , under the state of uniform ignorance, is the *uniform* prior, $P_0^{c,u}$, which assigns the same probability to every $H_i \in \mathcal{H}$, that is

$$P_0^{c,u}(H_i) = \frac{1}{n}, \quad i = 1, 2, \dots, n. \quad (7)$$

This prior probability function is, according to the objective Bayesian, the only function that DM is entitled to have, given her specific epistemic state of uniform ignorance⁵. The justification for this thesis runs as follows: Since DM possesses no specific probabilistic information, she is not in an epistemic position to favor one hypothesis over any other. Put differently, in order for DM to assign a different probability in H_i than in H_j , $i \neq j$, she must possess some evidence, upon which her probability differential is grounded. As Norton (2006) remarks: "...beliefs must be grounded in reasons, so that when there are no differences in reasons there should be no differences in beliefs" (2006, pp. 3-4). However, DM is assumed (counterfactually) to possess no such evidence. Hence, (7) seems to be the only sensible way to form her prior credences in H_i , $i = 1, 2, \dots, n$.⁶ Then, DM's unconditional subjective probability of (every) $A \in \mathcal{F}$ can be

⁵Despite its intuitive appeal, POI is a highly controversial principle (see, for example, Novack 2010 for a clear exposition of the sources of this controversy).

⁶More formally, $P_0^{c,u}$ is the (unique) credence function in $\mathcal{C}(\mathcal{H})$ that maximizes *Shannon's*

calculated according to the following equation, implied by the "theorem of total probability"⁷:

$$P_0^{c,u}(A) = \sum_{i=1}^n P_0^{c,u}(A | H_i) P_0^{c,u}(H_i). \quad (8)$$

Furthermore, her probability of $A \in \mathcal{F}$ conditional on any piece of information I_k , $k = 1, 2, \dots$ can also be calculated according to

$$P_0^{c,u}(A | I_k) = \sum_{i=1}^n P_0^{c,u}(A | H_i \wedge I_k) P_0^{c,u}(H_i | I_k). \quad (9)$$

Once DM has concluded the task of calculating her unconditional and conditional probabilities for every $A \in \mathcal{F}$ and for every possible information item I_k that may encounter in the future, she can exit the counterfactual mode of thinking and mentally return to the "actual" time period $t = 1$.

Step III: Now that the prior credal house of DM has been built, it is time for her to process the information that she actually possesses at $t = 1$, i.e. I_1 , thus forming her posterior probability function P_1 . Bayesian rationality requires that this information be processed only through BC. Hence, for every $A \in \mathcal{F}$,

$$P_1(A) = P_0^{c,u}(A | I_1). \quad (10)$$

It is easy to show that P_1 is coherent (i.e. it obeys the rules of probability calculus) not only for the case in which P_0^c is the uniform distribution but for any other coherent prior probability function P_0^c (see Appendix for a formal proof of this claim for the case of Ellsberg's "three-colors-one-urn"). Let us emphasize that (10) must be interpreted as a relation expressing a "double commitment" on the part of DM: (i) DM is committed to adopt $P_0^{c,u}(A | I_1)$ (formed in the epistemic state in which I_1 was a contingency) as her current credence $P_1(A)$ of A . This commitment expresses DM's willingness to be "dynamically consistent" with respect to her beliefs. (ii) DM is committed to assign the same degree of epistemic reliability to all posterior probabilities $P_1(A)$, $A \in \mathcal{F}$ formed via (10).

In the case of Ellsberg's "three-color-one-urn", the aforementioned Bayesian training process may be implemented as follows: DM (being at $t = 1$) must initially ignore the information I_1 : "There are 30 red balls in the urn". At this counterfactual epistemic state of no-information, she has to contemplate (among other probabilistic assignments) her subjective probability of "red" (R), conditional on I_1 as well as the corresponding conditional probabilities of "black" (B) and "yellow" (Y). Assuming that DM adopts POI, these subjective probabilities, denoted by $P_0^{c,u}(R | I_1)$, $P_0^{c,u}(B | I_1)$ and $P_0^{c,u}(Y | I_1)$ respectively can be

entropy SE (see Shannon 1948, Jaynes 1957). Maximization of SE (MaxEnt) is an epistemic principle which states that the prior that best represents DM's state of knowledge is the one that maximizes SE. Hence, if DM views MaxEnt as a normatively appealing principle, then her adoption of $P_0^{c,u}$ follows logically. Any prior probability function other than $P_0^{c,u}$ implies information that DM does not have.

⁷Of course the theorem of total probability continues to apply when the uniform prior $P_0^{c,u}$ is replaced by any other coherent probability function P_0^c .

calculated by means of (9). It follows that

$$\begin{aligned} P_0^{c,u}(R \mid I_1) &= \frac{1}{3} \\ P_0^{c,u}(B \mid I_1) &= \frac{1}{3} \\ P_0^{c,u}(Y \mid I_1) &= \frac{1}{3}. \end{aligned}$$

Once this step is completed, then DM should conform to BC, thus forming the following "actual" subjective probabilities at $t = 1$:

$$\begin{aligned} P_1(R) &= \frac{1}{3} \\ P_1(B) &= \frac{1}{3} \\ P_1(Y) &= \frac{1}{3}. \end{aligned}$$

Moreover, according to the aforementioned "double commitment" DM should assign the same degree of epistemic reliability to $P_1(R)$, $P_1(B)$ and $P_1(Y)$. As a result, she will be indifferent between "betting on red" and "betting on black", thus exhibiting no AA. The obvious question is, of course, whether DM is convinced by BT's aforementioned arguments or maintains her initial disposition towards avoiding "black". More specifically, assume that DM finds BTP overall unconvincing. The next question she will be asked to answer is which step of BTP she disagrees with. Does she disagree about the possibility of eliciting her prior via counterfactual reasoning, i.e. disagrees with the first two steps of the BTP? Or does she accept these two steps but disagree with the third, namely the Bayesian conditionalization? If her disagreement is about the feasibility (or desirability) of counterfactual thinking in general, then the Bayesian trainer should provide additional argument for this type of thinking (see next subsection). If, however, DM has accepted the first two steps of BTP as reasonable but denies the third one, then she may be accused of being "dynamically inconsistent". Specifically, assume first that in the epistemic context \mathcal{E}_0 , DM has expressed the view that the probability of A given the information I_1 , is equal to p , i.e. she has stated that $P_0^c(A \mid I_1) = p$, with I_1 being a contingency. Assume further that when DM is transferred to the epistemic framework \mathcal{E}_1 (in which I_1 is a certainty), she sets her probability of A not equal to p (as she committed to do in the context of \mathcal{E}_0) but to q , with $q \neq p$. This means that DM violates her own diachronic commitment about the probability of A and this is what makes her dynamically inconsistent. Dynamic inconsistency, however, is not consistent with rationality, since a dynamically inconsistent DM is susceptible to a dynamic Dutch Book (see, for example, Teller 1973, Lewis 1999). To put it succinctly, DM's refusal of the third step of BTP, conditional on her acceptance of the first two steps, should be interpreted as a clear sign of irrationality.

As mentioned above, the success of BTP depends on DM's willingness and ability to think counterfactually. If DM refuses to subscribe to this suggestion

or agrees to attempt it but without success, then DM fails to implement the first two steps of BTP thus resisting Bayesian training (at least the particular training implied by BTP). At this point, BT may be accused of trying to correct a real but undesirable feature of DM (namely AA) by asking her to exhibit a desirable but unattainable attitude (i.e. counterfactual thinking). Therefore, DM may ask BT to provide further reasons to defend the proposed counterfactual thinking. Next subsection discusses this issue.

4.1 The Plausibility and Merits of "Counterfactual Thinking"

Are there any additional reasons - beyond the prevention of comparative ignorance - that would recommend the implementation of the counterfactual strategy BTP? As mentioned above, a sceptic DM is likely to require BT to provide some additional arguments for "counterfactual thinking" before being persuaded to accept the first two steps of BTP. To this end, the Bayesian trainer may provide the following arguments:

4.1.1 Instinctive versus Reasoned Beliefs

It must be recognized at the outset that a *rational* DM is not a mere, passive bearer of beliefs. Instead, she is an active builder of beliefs. In other words, a rational DM does not just have beliefs; she rather forms beliefs. This raises the question: How should DM's beliefs be formed in order to be rational? In the case that DM possesses a certain amount of probabilistic information, I_1 , the aforementioned question reduces to the following one: how should DM process this information in order to end up with rational beliefs? Gilboa, Postlewaite and Schmeidler (2012) make an interesting distinction between *raw* versus *reasoned preferences*. Raw preferences refer to DM's "instinctive tendency to prefer one alternative over another". On the other hand, reasoned preferences are based on a process of reasoning: "Roughly, the decision maker exhibits raw preferences if she first acts, and then possibly observes her own act and stops to think about it. The decision maker is involved in reasoned choice if she first thinks, then decides how to act." (2012, p. 21). In a similar fashion, we distinguish between *instinctive* versus *reasoned beliefs* corresponding to reactive versus reflective processing, respectively of the available probabilistic information. More specifically, DM's instinctive reaction to information generates a credal state that reflects, in Carnap's (1961) terminology, DM's *momentary inclinations to believe*. On the other hand, careful analysis of the existing information is likely to enable DM to find her Carnapian *permanent dispositions to believe*. To highlight the difference between reactive versus reflective processing of information consider the following example: Imagine a person, X, who at $t = 1$, obtains the sour information, I_1 : "X suffers from lung cancer". X wants to assess her probability of the proposition A : "X will live for another five years". This can be done in two alternative ways: (a) DM allows I_1 to affect her subjective probability of A in a reactive way. This means that she forms the subjective

probability $P_1^{I_1}(A)$ under the psychological influence of knowing that I_1 is actually true, namely $P_1^{I_1}(I_1) = 1$. As a result, she over-estimates her chances of survival by coming up with (say) $P_1^{I_1}(A) = 0.5$. (b) DM processes I_1 with serious deliberation. This involves answering the following counterfactual question: "What would my credence of A be, had I_1 been not known, but rather it were one of many alternative eventualities, yet to actualize?". In this case, DM tries to evaluate the counterfactual conditional credence $P_0^c(A | I_1)$, with $0 < P_0^c(I_1) < 1$. In more concrete terms, DM attempts to access the credence that she would have assigned to A , if she performed this task in a psychologically neutral state in which she did not actually know that she suffers from lung cancer. Instead of the question "what is my probability of A given *I am certain* that I actually suffer from lung cancer?", DM answers the question "what is my probability of A *were I to become certain* that I suffer from lung cancer?". In this psychological state, DM is likely to be more unbiased (given that her life faces a hypothetical rather than an actual danger), thus identifying her true belief of A , (say) $P_0^c(A | I_1) = 0.1$.

Apart from Carnap, there are many other philosophers who have suggested DM's endorsement of an initial subjective probability function, that has to be formulated without the direct influence of any specific information I_1 , even if DM actually knows I_1 , (see Lewis 1980, Levi 1980, Skyrms 1983, and more recently Meacham 2008 and Titelbaum 2013). Howson (1991) in particular, strongly recommends the deletion of any specific information I_1 from the body of information, upon which DM's prior probability is based, as an answer to "the problem of old evidence". The latter was introduced by Glymour (1980) and Gardner (1982) and was initially interpreted as a serious pitfall of BCT. In line with Carnap, Howson argues that in order for BCT to get off the ground, "the dispositional properties of the agent's belief structure", as reflected in P_0^c , must be identified (Howson 1984, pp. 246).

Another example of how DM's actual encounter with I_1 might affect her ability to judge her own probabilities objectively is offered by Gul and Pesendorfer (2001): "Consider an individual who must decide what to eat for lunch. She may choose a vegetarian dish or a hamburger. In the morning, when no hunger is felt, she prefers the healthy, vegetarian dish. At lunchtime, the hungry individual experiences a craving for the hamburger." Hence, DM faces a "conflict between her ex ante ranking of options and her short-run cravings" (2001, pp. 1403).

Another argument in favor of counterfactual probabilistic reasoning comes from Greek mythology.⁸ Ulysses knows already from $t = 0$ that when he will listen to sirens' song at $t = 1$, he will be so enchanted by it that he will underestimate the probability of suffering a lethal encounter with them. In an attempt to secure that at $t = 1$ he will not succumb to siren's temptation, but instead he will act according to his emotionally neutral probabilistic beliefs, made at $t = 0$, the Greek hero asked his comrades to tie him up to the mast of his ship.

⁸This example is usually referred to the philosophical literature as the problem of "Ulysses and the Sirens" (see, for example, Elster 1979).

These examples may be thought of as a special case of a more general phenomenon pertaining to how emotional distortions impair DM's overall ability to think objectively. Indeed, there is a plethora of empirical studies that document a negative relationship between DM's level of anxiety (which in our case is caused by DM's perception of I_1 as non-contingent) and her ability to perform abstract reasoning tasks (see, for example, Leon and Revelle 1985). On another interpretation, the psychological effect of I_1 may be thought of as a "situational moderator", which negatively affects DM's information processing skills (see Humphreys and Revelle 1984). A common implication of both interpretations is the following: if DM treats I_1 as certain (that is when $P_1(I_1) = 1$), then she may experience emotional biases, which in turn impair her ability to uncover her genuine probabilistic dispositions.

4.1.2 The Pervasiveness of Counterfactual Thinking

Another argument that BT can invoke in defense of counterfactual thinking is that the latter is a fairly widespread mode of thinking to which DM resorts (albeit unconsciously) many times on a daily basis. Mandel et. al. (2007) introduce counterfactual thinking as follows: "It is human nature to wonder how things might have turned out differently - either for the better or for the worse." Markman et. al. (2012) hold that "counterfactual thinking - the capacity to reflect on what would, could or should have been if events had transpired differently - is a pervasive, yet seemingly paradoxical human tendency." (2012, p. 175). Similarly, Kahneman and Miller (1986) argue that, quite often, a stimulus that DM experiences in the course of her actual life "selectively recruits its own alternatives ... and is interpreted in a rich context of remembered and constructed representations of *what it could have been, might have been, or should have been.*" (1986, p. 136, emphasis added). This means that certain events trigger counterfactual thinking in the context of which "events are sometimes compared to counterfactual alternatives that are constructed ad hoc rather than retrieved from past experience." (1986 p. 137). According to the so-called "simulation heuristics" (Kahneman and Tversky 1982, Tversky and Kahneman 1973), DM employs counterfactual (i.e. contrary-to-facts) thinking quite frequently in order to spontaneously construct "alternatives to experience". Comparing these counterfactual alternatives to actual experience can elicit either positive or negative psychological effects. More specifically, an *upward counterfactual*, i.e. a constructed alternative that improves on reality, produces a negative effect (usually a feeling of regret), whereas a *downward counterfactual*, i.e. an alternative that is worse than reality, activates a positive effect (see Roesse 1994). More generally, over the last four decades, psychologists have attempted to analyze the causes behind people's propensity to generate counterfactuals as well as the effects that this mode of thinking produces (see, for example, Mandel et. al. 2007).

In addition, BT can argue that counterfactual thinking is intimately related to causal thinking (see, for example, Spellman and Mandel 1999). Kment (2020) introduces the relation between counterfactuals and causality as follows:

"Counterfactual thought is an important element of our cognitive lives. In making practical decision, we are often led to *ask what would happen if we were to carry out a certain action*, and we frequently support causal claims by showing that the putative *effect depends counterfactually on the supposed cause*." (2020, p. 1 emphasis added). More specifically, in the context of Lewis's (1973) theory, causation is reduced to counterfactual dependence: "event c causes event e " if in the absence of c e would have not occurred: "We think of a cause as something that makes a difference, and the difference it makes must be a difference from *what would have happened without it*. Had it been absent, its effects - some of them, at least, and usually all - would have been absent as well." (1973, p. 161, emphasis added).⁹

For all the above reasons, counterfactual thinking is not a counter-intuitive form of thinking to which DM feels alien. On the contrary, the typical DM exhibits a tendency to raise "what-if questions" on many occasions during her everyday life. For example, imagine a DM who arrived late to her morning meeting due to heavy traffic. DM is likely to ask (consciously or subconsciously) the question "what would happen if the traffic were not so heavy" and give the answer "if the traffic were not so heavy, I would be on time for my meeting". This type of thinking is what is required of DM in order to successfully implement the first two steps of BTP. Given the importance and pervasiveness of counterfactual thinking for DM's cognitive life, any objection by DM to applying the first two steps of BTP seems (at least in BT's eyes) unwarranted.

5 Conclusions

The main points of the paper are summarized in the form of the following hypothetical dialogue between a decision maker (DM) and a Bayesian trainer (BT). DM is supposed to know the decision problem defined by Ellsberg in the context of his "three-colors-one-urn" thought experiment.

DM: When asked what color I would like to bet on in the context of the Ellsberg "three-colors-one-urn" case, I answer without hesitation "red".

BT: Why not "black"? After all, this color also ensures the same monetary outcome of 100 euros.

DM: "Black" can ensure the same monetary result, but the objective probability of "black" is unknown to me while the probability of "red" is known and equal to $1/3$.

BT: So am I right to conclude that the subjective probability you assign to "red" is greater than the one you assign to "black"?

DM: Not necessarily, what I told you is that the probability of "red" is known to me. I didn't state that the probability of "red" is necessarily greater than that of "black".

⁹Since 1973, the year Lewis introduced his Counterfactual Analysis of Causation, there has been intense debate among philosophers about whether causality can be analyzed in terms of counterfactual dependence (see the entry "Counterfactual Theories of Causation" in Stanford Encyclopedia of Philosophy.)

BT: Indeed, you did not make that statement explicitly. However, I was able to deduce this implicitly by analyzing your betting behavior in conjunction with the assumption that you adhere to the SEUM criterion of choice. Am I right to make this assumption? Do you find this criterion rational?

DM: Yes. Intuitively, I think that between two actions, it makes sense to choose the one with the highest expected utility.

BT: In that case, you have a problem. You can't have your cake and eat it too. Your acceptance of the SEUM criterion together with your preference to bet on "red" than on "black" and the fact that the monetary outcomes are identical for both bets entail the conclusion *C*: "Your subjective probability of "red" is larger than that of "black". (At this point BT shows analytically to DM how the aforementioned three premises imply *C*).

DM: I see your point. I cannot prefer 'red' to 'black', accept the SEUM criterion of choice, while at the same time claiming that I do not necessarily consider 'red' more probable than 'black'. Nonetheless, I still feel that what pushes me towards "red" is that I consider it more reliable (rather than more probable) than "black".

BT: This means that you allow a new factor, namely the "degree of reliability" of your probabilistic assessments, to enter the decision-making process. This of course means that you have effectively rejected the SEUM criterion, as the latter is based entirely on your subjective probabilities and utilities and not at all on factors such as the reliability of your subjective probabilities.

DM: What a predicament! If I accept the validity of the SEUM criterion, I must admit that I find "red" more likely than "black". On the other hand, if I insist that my subjective probability of "red" is not necessarily greater than that of "black," then I should deny the validity of the SEUM criterion. Is there a way out of this cognitive dissonance?

BT: I shall try to help you. First let me ask you, what is your subjective probability of "black"?

DM: I am not sure. It depends on how many black balls are in the urn. Since I know that 30 of the 90 balls are red, I conclude that the number of black balls is between 1 and 59.

BT: Yes, but if you were forced to bet on black, what would your odds be? I want you to try and elicit a sharp and unique number as your subjective probability of "black".

DM: Let me think. I know that the objective probability of "either black or yellow" is $2/3$. Moreover, I do not have any information pertaining to the relative frequency of black or yellow balls in the urn. Hence, I find it reasonable to divide $2/3$ by 2, thereby concluding that my subjective probability of "black" is $1/3$.

BT: So your subjective probability of "black" is equal to that of "red". Yet, you still prefer to bet on "red" rather than "black".

DM: Yes, because as I told you before it's not that I consider "red" more likely than "black". I assign both colors the same probability of $1/3$. However, I consider $1/3$ of "red" more reliable than $1/3$ of "black". I repeat, this is the reason that urges me to avoid "black" and prefer "red".

BT: Now we seem to be getting somewhere. The next question I want you to think about is why you find the probability of "red" more reliable than the (equal) probability of "black".

DM: I don't have to think much to answer this question. I consider the probability of "red" to be more reliable than that of "black" because I have been given information about the objective probability of "red" (in the form of the relative frequency of red balls in the urn) while I possess no such information for the case of "black".

BT: So what you claim is that the probabilistic information about "red" makes you feel epistemically superior to "red", or equivalently epistemically inferior to "black". In other words, you feel comparatively more ignorant about "black" than "red". Am I correct in inferring that this comparative ignorance is the main reason behind your aversion to bet on "black"?

DM: Yes, I think what you just said pretty much explains my behavior.

BT: So, if the cause of ambiguity aversion is the epistemic state of comparative ignorance, then if somehow you managed to get out of this state and into the state of uniform ignorance (in which you are equally uninformed about the three colors in the urn), you would cease to exhibit an aversion towards "black".

DM: Can you elaborate on this point a bit more?

BT: Suppose you are faced with a situation in which you have no information about the relative proportions of the three colors in the urn. Then, do you find it reasonable to assign equal probabilities to each color, i.e. $1/3$?

DM: Yes.

BT: In this case, you consider the so-called "Principle of Indifference" to be a sound epistemic principle. This principle says that any difference in your probabilistic assignments of the three colors should be accounted for by a corresponding difference in your probabilistic information about those colors. Since you have identical information about the three colors, (i.e. zero) you should also have identical subjective probabilities for the three colors (i.e. $1/3$).

DM: Yes, but the situation I'm dealing with is different. In my case, I have information about the relative frequency of red balls.

BT: Indeed you have. But we'll get to that in a minute. For now, I want you to stay in the hypothetical epistemic state of zero information and tell me whether, in that state, your probabilistic evaluations are characterized by the same degree of epistemic reliability. In other words, is the $1/3$ you assign to "red" just as reliable as the $1/3$ you assign to "black" or "yellow"?

DM: Yes it is.

BT: So, am I right to assume that you are indifferent between betting on "red" and betting on "black"?

DM: Yes, you are.

BT: Now, allow me to take the discussion a step further and ask you about your "conditional probabilities" while still remaining in the state of zero information. More specifically, what is your probability of "red" conditional on the information that there are 30 red balls in the urn?

DM: $1/3$.

BT: And what is your probability of "red" conditional on the information that there are 40 red balls in the urn?

DM: 4/9.

BT: Now a slightly different question: what is your probability of "black" conditional on the information that there are 40 red balls in the urn?

DM: It is 5/9 divided by 2, that is 5/18.

BT: Here comes a crucial question. Conditional on the information that there are 40 red balls in the urn, do you find the above conditional probability of "red" more reliable than the corresponding conditional probability of "black"? In other words, do you find 4/9 more reliable than 5/18?

DM: Not really. I see the whole process as a probabilistic exercise, in which I analyze how my previous probabilistic assignments change in light of various hypothetical information scenarios. If I ever receive the aforementioned information that there are 40 red balls in the urn, then my prior probability of "red" changes from 1/3 to 4/9, and my prior probability of "black" (and "yellow") changes from 1/3 to 5/18. Alternatively, if I receive another piece of information, for example, there are 10 black balls in the urn, then my prior probability of "black" changes from 1/3 to 1/9, while the prior probability of "red" changes from 1/3 to 4/9. All these probabilistic assignments refer to hypothetical epistemic states, which means that they are all "exercises on paper" and therefore psychologically neutral and "equally reliable".

BT: What I understand from what you have told me is that the information "there are 30 red balls in the urn" has different psychological effects on you depending on its modal status. More specifically, if the information is certain, then it puts you in a state of comparative ignorance. Conversely, if the information is contingent - on par with any other possible information you might receive in the future - then it produces no such results.

DM: I guess you can put it this way. Taking a piece of information as certain is quite different from taking it as merely possible.

BT: Let us now take our discussion one step further. I want you to visualize two distinct points in time, namely $t = 0$ and $t = 1$. At $t = 0$, you possess no probabilistic information about the relative proportions of the three colors in the urn. Being at this epistemic (information-free) state, you build your prior subjective probabilities (both unconditional and conditional ones) for all relevant events of interest. For example, you decide (as you did before) that your unconditional probability of "red" is 1/3, your probability of "red" conditional on the information that there are 40 red balls in the urn is 4/9, your unconditional probability of "black" is 1/3, your probability of "black" conditional on the information that there are 50 black balls in the urn is 5/9 and so on. Moreover, still at $t = 0$, commit to the following: a) If at the future time period $t = 1$, you receive the probabilistic information I_1 (e.g. "there are 40 red balls in the urn") then you will set your posterior probability of "red" at $t = 1$ equal to your prior probability of "red" conditional on I_1 formed at $t = 0$. Moreover, this commitment of yours covers all your probabilistic evaluations at $t = 1$. For example, your posterior probability of "black" at $t = 1$ should be set equal to your prior probability of "black" conditional on I_1 and so on. b) You will treat

all these posterior probabilities set at $t = 1$ as equally reliable. For example, your posterior probability of "red" enjoys the same degree of epistemic reliability as your posterior probability of "black". Do you think that these two commitments are reasonable? Are you prepared to accept them as canons of rationality?

DM: I think I am. The way you put the problem leaves me little room for disagreement. If I reject the two commitments you mentioned, then it is like expressing an opinion at $t = 0$ and then without any reason changing that opinion at $t = 1$. In other words I would appear to be "dynamically inconsistent".

BT: In this case at $t = 1$, you will end up with a Bayesian subjective probability function. Your probabilistic assessments combined with the principle of expected utility maximization will produce choices immune to ambiguity aversion.

DM: Yes, you are probably right as long as you are dealing with a decision-making problem in which there are the two distinct time periods you mentioned above, namely $t = 0$ and $t = 1$. In the context of the Ellsberg "three-colors-one-urn" case with which we began our discussion, however, the first time period $t = 0$ is missing. At the moment of my decision between "betting on red" and "betting on black" I already possess the information "there are 30 red balls in the urn". In other words, there is no privileged, information-free time point $t = 0$, at which I would be at the ideal psychological state to build my prior probability function.

BT: In this case why don't you create the elusive time point $t = 0$ artificially or *counterfactually*. In other words, why don't you "pretend" not to know the information I_1 : "there are 30 red balls in the urn" and proceed to construct your prior probability function in the same way as you would in the case that you actually do not know I_1 ?

DM: You mean to perform a "mental simulation" in the context of which I reproduce (counterfactually) the neutral epistemic state in which I possess no probabilistic information? I am not sure that I can do that.

BT: I know it is a rather difficult undertaking. On the other hand, you enter a counterfactual mode of thinking many times on a daily basis. In any case, your conversion from an ambiguity-averse decision maker to a Bayesian one hinges precisely upon the successful implementation of this task.

References

- Al-Najjar, N. I., and Weinstein, J. (2009). The ambiguity aversion literature: a critical assessment. *Economics & Philosophy*, 25(3), 249-284.
- Arlo-Costa, H., & Helzner, J. (2005). Comparative ignorance and the Ellsberg phenomenon.
- Becker, S. W., and Brownson, F. O. (1964). What price ambiguity? Or the role of ambiguity in decision-making. *Journal of political economy*, 72(1), 62-73.
- Carnap, R. (1962). *Logical Foundations of Probability*. 2nd ed. University of Chicago Press.
- Carnap, R. (1971). Inductive logic and rational decision. *Studies in Inductive Logic and Probability*, 1, 5-31.
- Chow, C. C., and Sarin, R. K. (2001). Comparative ignorance and the Ellsberg paradox. *Journal of risk and Uncertainty*, 22, 129-139.
- Curley, S.P., Yates, J.F. and Abrams, R.A. (1986). Psychological sources of ambiguity avoidance. *Organizational behavior and human decision processes*, 38(2), 230-256.
- Einhorn, H. J., and Hogarth, R. M. (1985). Ambiguity and uncertainty in probabilistic inference. *Psychological review*, 92(4), 433.
- Einhorn, H. J., and Hogarth, R. M. (1986). Judging probable cause. *Psychological bulletin*, 99(1), 3.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 75, 643-669.
- Elster, J. (1979). *Ulysses and the sirens: Studies in rationality and irrationality*.
- Fellner, W. (1961). Distortion of subjective probabilities as a reaction to uncertainty. *The Quarterly Journal of Economics*, 75(4), 670-689.
- Fox, C. R. and Tversky, A. (1995). Ambiguity aversion and comparative ignorance. *The Quarterly Journal of Economics*, 110(3), 585-603.
- Fox, C. R. and Weber, M. (2002). Ambiguity aversion, comparative ignorance, and decision context. *Organizational Behavior and Human Decision Processes*, 88, 476-498.
- Frisch, D. and Baron, J. (1988). Ambiguity and rationality. *Journal of Behavioral Decision Making*. 1:149-157.
- Gardner, M. R. (1982). Predicting novel facts. *The British Journal for the Philosophy of Science*, 33(1), 1-15.

- Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with a nonunique prior. *Journal of Mathematical Economics*, 18, 141–153.
- Gilboa, I. and Marinacci, M. (2013). Ambiguity and the Bayesian Paradigm. In *Advances in Economics and Econometrics*, Volume 1, Economic Theory: Tenth World Congress (Vol. 49, p. 179). Cambridge University Press.
- Gilboa, I. and Marinacci, M. (2016). Ambiguity and the Bayesian paradigm. In *Readings in formal epistemology*, 385-439. Cham: Springer International Publishing.
- Gilboa, I., Postlewaite, A., and Schmeidler, D. (2012). Rationality of belief or: why Savage's axioms are neither necessary nor sufficient for rationality. *Synthese*, 187, 11-31.
- Glymour, C. (1980). *Theory and Evidence*, Princeton: Princeton University Press.
- Gul, F. and Pesendorfer, W. (2001). Temptation and self-control. *Econometrica*, 69(6), 1403-1435.
- Heath, C. and Tversky, A. (1991). Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty*, 4(1), 5-28.
- Hogarth, R. M., & Kunreuther, H. (1985). Ambiguity and insurance decisions. *The American Economic Review*, 75(2), 386-390.
- Howson, C. (1984). Bayesianism and support by novel facts. *The British Journal for the Philosophy of Science*, 35(3), 245-251.
- Howson, C. (1995). Theories of probability. *The British Journal for the Philosophy of Science*, 46(1), 1-32.
- Humphreys, M. S. and Revelle, W. (1984). Personality, motivation, and performance: a theory of the relationship between individual differences and information processing. *Psychological review*, 91(2), 153.
- Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical review*, 106(4), 620.
- Kahneman, D. and Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological review*, 93(2), 136.
- Kahneman, D. and Tversky, A. (1982). The psychology of preferences. *Scientific American*, 246(1), 160-173.
- Kment, B. (2020). Counterfactuals and causal reasoning. In *Perspectives on Causation: Selected Papers from the Jerusalem 2017 Workshop* (pp. 463-482). Springer International Publishing.

- Langer, E. J. (1975). The illusion of control. *Journal of personality and social psychology*, 32(2), 311.
- Leon, M. R. and Revelle, W. (1985). Effects of anxiety on analogical reasoning: a test of three theoretical models. *Journal of Personality and Social Psychology*, 49(5), 1302-1315.
- Levi, I. (1980). The enterprise of knowledge: An essay on knowledge, credal probability, and chance. MIT press.
- Levi, I. (1985). Imprecision and indeterminacy in probability judgment. *Philosophy of Science*, 52, 390-409.
- Lewis, D. (1973). Counterfactuals and comparative possibility. In IFS: Conditionals, Belief, Decision, Chance and Time (pp. 57-85). Dordrecht: Springer Netherlands.
- Lewis, D. (1980). A subjectivist's guide to objective chance. In IFS: Conditionals, Belief, Decision, Chance and Time (pp. 267-297). Dordrecht: Springer Netherlands.
- Lewis, D. (1999). Papers in Metaphysics and Epistemology: Volume 2. Cambridge University Press.
- MacCrimmon, K. R. and Larson, S. (1979). Utility theory: axioms versus paradoxes. *Expected Utility Hypotheses and the Allais Paradox* (ed. Allais, M. and Hagen, O.), 333-410.
- Machina, M. J., & Siniscalchi, M. (2014). Ambiguity and ambiguity aversion. In *Handbook of the Economics of Risk and Uncertainty*, Vol. 1, pp. 729-807.
- Mandel, D. R., Hilton, D. J. and Catellani, P. (2007). The psychology of counterfactual thinking. Routledge.
- Markman, K. D., Karadogan, F., Lindberg, M. J. and Zell, E. (2012). Counterfactual thinking: Function and dysfunction. In *Handbook of imagination and mental simulation* (pp. 175-194). Psychology Press.
- Meacham, C. J. (2008). Sleeping beauty and the dynamics of de se beliefs. *Philosophical Studies*, 138, 245-269.
- Meacham, C. J. (2016). Ur-priors, conditionalization, and ur-prior conditionalization. *Ergo: An Open Access Journal of Philosophy*, 3.
- Muthukrishnan, A.V., Wathieu, L. and Jing Xu, A. (2009). Ambiguity aversion and the preference for established brands. *Management Science*, 55, 1933-1941.
- Norton, J. D. (2008). Ignorance and indifference. *Philosophy of Science*, 75(1), 45-68.

- Raiffa, H. (1961). Risk, Ambiguity, and the Savage Axioms: Comment. *The Quarterly Journal of Economics*, 75(4), 690-694.
- Roberts, H. V. (1963). Risk, ambiguity, and the Savage axioms: Comment. *The Quarterly Journal of Economics*, 77(2), 327-336.
- Roese, N. J. (1994). The functional basis of counterfactual thinking. *Journal of personality and Social Psychology*, 66(5), 805.
- Rubaltelli, E., Rumiati, R., and Slovic, P. (2010). Do ambiguity avoidance and the comparative ignorance hypothesis depend on people's affective reactions?. *Journal of Risk and Uncertainty*, 40, 243-254.
- Sahlin, N. E. (1993). Notes on Philosophy, Probability and Mathematics.
- Savage, L. (1954). The foundations of statistics. New-York, John Wiley.
- Savage, L. (1962). Bayesian statistics. In *Recent developments in decision and information processes* (R. E. Machol and P. Gray, eds.) 161-194, Macmillan New York.
- Savage, L. (1972). The foundations of statistics. Courier Corporation.
- Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica*, 57(3), 571-587.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell system technical journal*, 27(3), 379-423.
- Shapiro, D. (2020). Limited Impact of Business Development Programs on Profitability in the Presence of Ambiguity Aversion. Asian Development Bank Economics Working Paper Series, (614).
- Skyrms, B. (1983). Three ways to give a probability assignment a memory.
- Slovic, P. and Tversky, A. (1974). Who accepts Savage's axiom?. *Behavioral science*, 19(6), 368-373.
- Spellman, B. A. and Mandel, D. R. (1999). When possibility informs reality: Counterfactual thinking as a cue to causality. *Current Directions in Psychological Science*, 8(4), 120-123.
- Strevens, M. (2006). The Bayesian approach to the philosophy of science. For the Macmillan Encyclopedia of Philosophy, 2nd edition.
- Teller, P. (1973). Conditionalization and Observation. *Synthese*, 26: 218-258.
- Titelbaum, M. G. (2013). Quitting certainties: A Bayesian framework modeling degrees of belief. Oxford University Press.
- Toda, M., and Shuford Jr, E. H. (1965). Utility, induced utilities, and small worlds. *Behavioral Science*, 10(3), 238-254.

- Tversky, A. and Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive psychology*, 5(2), 207-232.
- Yates, J. F., and Zukowski, L. G. (1975). The anatomy and consequences of ambiguity in decision making. University of Michigan, Department of psychology.
- Walley, P. (1991). Statistical reasoning with imprecise probabilities (Vol. 42). London: Chapman and Hall.