

7-2005

Comparative Ignorance and the Ellsberg Phenomenon

Horacio Arlo-Costa

Carnegie Mellon University, hcosta@andrew.cmu.edu

Jeffrey Helzner

Columbia University, jh2239@columbia.edu

Follow this and additional works at: <http://repository.cmu.edu/philosophy>

 Part of the [Philosophy Commons](#)

Published In

.

This Conference Proceeding is brought to you for free and open access by the Dietrich College of Humanities and Social Sciences at Research Showcase @ CMU. It has been accepted for inclusion in Department of Philosophy by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Comparative Ignorance and the Ellsberg Phenomenon

Horacio Arlo-Costa
Carnegie Mellon University
hcosta@andrew.cmu.edu

Jeffrey Helzner
Columbia University
jh2239@columbia.edu

Abstract

The “Ellsberg phenomenon” has played a significant role in research on imprecise probabilities. Fox and Tversky [7] have attempted to explain this phenomenon in terms of their “comparative ignorance” hypothesis. We challenge that explanation and suggest that our recent empirical work suggests an explanation that is much closer to Ellsberg’s own diagnosis.

Keywords. Ellsberg, comparative ignorance, ambiguity aversion

1 Introduction

In “decision making under risk” the decision maker has access to an objective probability over the states of nature. In “decision making under uncertainty” the decision maker does not have access to such information. The decision theoretic significance of this distinction is denied by those subjectivists who maintain that the recommendation to maximize expected utility is as applicable to uncertainty as it is to risk. When confronted with decision making under uncertainty, the rational decision maker does have a subjective probability distribution over the states of nature. In such a situation, the rational decision maker ought to select an alternative that maximizes subjective expected utility. It is in this sense that certain subjectivists (e.g. Savage) regard decision making under uncertainty as reducible to decision making under risk.

The appropriateness of this reduction was called into question by Ellsberg [4], [5]. Ellsberg offered several examples that he regarded as compelling evidence against the indicated reduction. Empirical work has demonstrated that Ellsberg’s examples do generate a significant number of responses that are incompatible with all theories that adopt the standard relationship between preference and choice and make the further

assumption that preference satisfies “weak ordering” and “the sure-thing principle”. In particular, such “deviant” responses are incompatible with Savage’s subjective expected utility theory [SEU].

Ellsberg’s work revived an earlier tradition first articulated in the writings of Knight [16] and Keynes [15] among others. Knight in particular distinguished between *risk*, which can be represented by precise probabilities and (*unmeasurable*) uncertainty, which does not admit such representations. Keynes, on the other hand, distinguished between judged *probability*, which reflects the proportion of evidence in favor of a proposition, and the *weight* of evidence, which measures the quantity of evidence supporting the judged proportion.

At least two distinct interpretations of the deviant responses to Ellsberg’s examples have been considered. One interpretation takes these responses as further evidence of our limitations in rational decision-making and of the need to discharge certain rationality assumptions in the foundations of economics. Such an interpretation is congenial to those research programs that are concerned to develop a descriptively adequate account of decision-making under uncertainty. While this first interpretation assumes that SEU (or something very similar) is the appropriate standard of rationality, there is a second interpretation that takes the deviant responses as evidence that the prevailing normative account is flawed (or at least incomplete).

Although this second interpretation is an anathema to some elements of the Bayesian orthodoxy, there is some prima facie justification for the interpretation. Among the deviant respondents are a significant number of sophisticated decision makers who maintain their choices even after they are reminded that their choices are incompatible with SEU. It should be noted that Ellsberg himself was concerned with this second interpretation of the deviant responses given to his now famous examples. Regarding these responses, Ellsberg wrote:

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 way, even though they may be generally respectful of the Savage axioms. There are strong indications, in other words, not merely of the existence of reliable patterns of blind behavior but of operations of definitive normative criteria, differing from and conflicting with the familiar ones, to which these people are trying to conform. (Ellsberg, [4])

Several normative theories are capable of accommodating the indicated responses. Notable examples include Ellsberg ([4], [5]), Levi ([17], [18]) and Gärdenfors and Sahlin ([9]). There are significant differences among some of these theories. For example, Ellsberg ([4], [5]) and Gärdenfors and Sahlin ([9]) maintain the weak ordering assumption and relax the sure-thing principle while Levi ([17], [18]) relaxes the weak ordering assumption and maintains the sure-thing principle. Despite these differences, one thing that is common to each of the theories above is the use of imprecise probabilities.

2 Ambiguity Aversion and Comparative Ignorance

In contrast to normative accounts such as those mentioned in the previous section, behavioral decision theorists have sought to explain observed deviations from SEU in terms of psychological effects. The deviant responses associated with Ellsberg’s examples have been attributed to a psychological effect known as “ambiguity aversion”. According to this usage of the term, ambiguity is present to the extent that the decision maker’s information is insufficient to establish a probability distribution over the relevant states of nature.

Results from several studies (Heath and Tversky [12], Cohen and Hansel [2], and Howell [13]) suggest that an unqualified ambiguity aversion explanation is untenable. In particular, Heath and Tversky [12] found that people are not averse to betting in the presence of ambiguity when they believe they that they are especially competent or knowledgeable. Such results prompted Fox and Tversky [7] to investigate the following explanation:

When evaluating an uncertain event in isolation people attempt to assess its likelihood – as a good Bayesian would – paying rel-

atively 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. The main implication of this account, called the *comparative ignorance hypothesis*, is that ambiguity aversion will be present when subjects evaluate clear and vague prospects jointly, but it will greatly diminish or disappear when they evaluate each prospect in isolation. (Fox and Tversky, [7])

Fox and Tversky conducted several studies to show that comparative ignorance drives ambiguity aversion. In one study from [7], 141 undergraduates at Stanford University were divided into three groups: Noncomparative Clear, Noncomparative Vague, and Comparative. The study used questions concerning draws from two bags of poker chips, Bag A and Bag B. Each of the bags contained exactly 100 chips. Chips were assumed to come in two colors, black and red. Bag A was known to contain exactly 50 red chips and 50 black chips. The proportions for Bag B were unknown. Subjects in Noncomparative Clear were asked to specify the most they would be willing to pay for a ticket that pays either \$100 or \$0 depending on whether a particular outcome obtains as a result of a random selection from Bag A. Subjects in Noncomparative Vague were asked to perform the same task with respect to Bag B. Subjects in Comparative were asked to price both tickets (i.e. they were asked to perform the described task with respect to each of the bags).

As predicted by the comparative ignorance hypothesis, the mean prices for Noncomparative Clear and Noncomparative Vague were relatively close together. The mean for Noncomparative Clear was \$17.94. The mean for Noncomparative Vague was \$18.42. As predicted by the comparative ignorance hypothesis, the corresponding means for Comparative were relatively far apart. Within the Comparative group, the mean for the “clear bet” (i.e. the bet on Bag A) was \$24.34 while the mean for the “vague bet” (i.e. the bet on Bag B) was \$14.85.

The distinction between comparative and non-comparative contexts is crucial in order to interpret Fox and Tversky’s results. In comparative contexts each participant evaluates some “clear” lotteries (i.e. lotteries having clear probabilities) and some “vague” lotteries (i.e. lotteries having vague probabilities). In non-comparative contexts different participants evaluate just one type lottery; that is, a participant in a

non-comparative context evaluates either clear lotteries or vague lotteries but not both. For Fox and Tversky the distinction between comparative and non-comparative *assessments* is ultimately subjective. So, they refuse to trace a clear distinction between comparative and non-comparative contexts only in terms of objective manipulations of the experimental set-up. What counts for them is the state of mind of the decision-maker. So, for example, Fox and Tversky remark:

[...] there is no guarantee that subjects in the comparative conditions actually perform the suggested comparison, or that subjects in the non-comparative conditions did not independently generate a comparison. In Ellsberg's two-color problem, for example, [some] people who are presented with the vague urn alone may spontaneously invoke a comparison to 50-50 urn, especially if they have previously encountered such a problem. (Fox and Tversky, [7])

This leads Fox and Weber [8] to remark that '[...] the presence of ambiguity aversion in non-comparative contexts does not necessarily contradict the comparative ignorance hypothesis.'

Moreover some recent studies, like the one conducted by Chow and Sarin in [1] show that subjects prefer clear (known) probabilities over vague (unknown) probabilities and that this preference does persist even in non-comparative conditions. Fox and Weber are aware of these findings but they do not think that they threaten the credibility of the comparative ignorance hypothesis. It is unclear though whether one can assume that in Chow and Sarin's experiment subjects perform psychological comparisons not elicited by manipulation of the experimental context. Although Fox and Tversky are right that the presence of *ambiguity aversion* in non-comparative settings does not contradict the comparative ignorance hypothesis, this does not mean that the non-Bayesian assessment of uncertain prospects in non-comparative settings can always be explained in terms of an aversion effect. Moreover the behavior of agents in comparative settings can also be better explained in terms of models that treat imprecise probabilities seriously.

Even more recent work [8] has tried to provide support for the comparative hypothesis in a way that does not relay on the comparative/non-comparative elicitation paradigm. We will comment on these findings below although most of our results (especially the ones that are incompatible with the comparative ignorance hypothesis) reproduce the experimental set-up utilized by Fox and Tversky (to elicit non-comparative condi-

tions).

Ultimately the central point that interest us is to what extent the postulation of psychological effects (like the comparative hypothesis) manages to produce (or not produce) comprehensive and unified explanations of Ellsberg's phenomenon. We do not doubt that an assortment of psychological effects might play a role in describing choice behavior under uncertainty. In the case of Ellsberg's phenomenon there are also other kinds of relevant effects which have been proposed recently (like the one based on avoiding a 'stacked deck' possibility, rather than avoiding ambiguity [3]). These phenomena might be salient in various experimental and real situations.

But these explanations seem to have a limited scope. At least the strategy of explaining vagueness away in terms of an alleged pathology in choice behavior does not seem to be powerful enough to explain the data that we are about to present. Relatively mild extensions of the Ellsberg's original experiment seem to be better explained in terms of Ellsberg's own normative ideas about agents paying serious attention to irreducible second-order characteristics such as vagueness, than in terms of psychological effects.

3 Our Study

Our study involved 97 undergraduates at Carnegie Mellon University. At the time that the study was conducted, each of the subjects was enrolled in 80-100, an introductory philosophy course that is taught at CMU. The subjects were divided into four groups: Clear I, Vague I, Clear II, and Vague II. Each recitation section for 80-100 was assigned to a group. Questionnaires were administered in recitation sections. We used four different questionnaires, one for each of the groups. Each questionnaire consisted of four modules. These modules were based on the following scheme:

Poker chips come in two colors: black and white. The following three questions are to be considered under the assumption that X.

Question 1 What is the largest amount of money you would be willing to pay for a ticket to play a game where:

You win \$55 if a black chip is drawn.
You win \$0 if a white chip is drawn.

Write your answer here:

Question 2 Would you be willing to pay \$Z (for some positive value of Z) to play a game where:

You win \$55 + \$20Z if a black chip is drawn. You win \$0 if a white chip is drawn.

Circle the appropriate answer below:

YES NO

If your answer to Question 2 is ‘Yes’, then continue to Question 3. Otherwise, skip Question 3 and continue on with the rest of the questionnaire.

Question 3 What is the largest positive amount \$Z that you would be willing to pay to play a game where:

You win \$55 + \$20Z if a black chip is drawn. You win \$0 if a white chip is drawn.

Write your answer here: Continue on with the rest of the questionnaire.

Different values of X in the opening paragraph lead to different modules. Assume that Bag A contain exactly 100 poker chips, 50 black chips and 50 white chips. Assume that Bag B contain exactly 100 poker chips but that the ratio of black chips to white chips is unknown. We considered the following values for the assumption X :

Clear $X =$ “the randomly selected chips will be taken from Bag A.”

Vague $X =$ “the randomly selected chips will be taken from Bag B.”

%20 Clear - %80 Vague $X =$ “there is a %20 chance that the randomly selected chip will be taken from Bag A and an %80 chance that the randomly selected chip will be taken from Bag B.”

%50 Clear - %50 Vague $X =$ “there is a %50 chance that the randomly selected chip will be taken from Bag A and a %50 chance that the randomly selected chip will be taken from Bag B.”

%80 Clear - %20 Vague $X =$ “there is an %80 chance that the randomly selected chip will be taken from Bag A and a %20 chance that the randomly selected chip will be taken from Bag B.”

We can now describe the questionnaires that were given to the two I-groups. The questionnaire for Clear I consisted of the following modules: Clear, %80 Clear - %20 Vague, %50 Clear - %50 Vague, and %20 Clear - %80 Vague. The questionnaire for Vague I consisted of the following modules: Vague, %20 Clear - %80 Vague, %50 Clear - %50 Vague, and %80 Clear - %20 Vague. In both cases the modules appearing in the questionnaire were arranged in the order stated above (i.e. decreasing clarity for Clear I and decreasing vagueness for Vague I).

The questionnaires for the II-groups involved two additional modules, Clear-2 and Vague-2. Clear-2 was obtained by removing Question 2 from Clear I. Vague-2 was obtained by removing Question 2 from Vague I.¹ We can now describe the questionnaires that were given to the two II-groups. The questionnaire for Clear II consisted of the following modules: Clear-2, %80 Clear - %20 Vague, %50 Clear - %50 Vague, and %20 Clear - %80 Vague. The questionnaire for Vague II consisted of the following modules: Vague-2, %20 Clear - %80 Vague, %50 Clear - %50 Vague, and %80 Clear - %20 Vague. In both cases the modules appearing in the questionnaire were arranged in the order stated above (i.e. decreasing clarity for Clear II and decreasing vagueness for Vague II).

Each cell in Table 1 contains the mean (m) and the standard deviation (s) for the responses to the first module of the questionnaire that was administered to the associated group. The comparative ignorance hypothesis predicts that the Question 1 means for Clear and Vague (both I and II) should be relatively close, but, as shown in Table 1, this is not the case. Similar remarks apply to the Question 3 means.

Group	Question 1	Question 3
Clear I	m = 15.33, s = 9.34	m = 44.80, s = 36.43
Vague I	m = 5.42, s = 9.23	m = 15.29, s = 12.32
Clear II	m = 13.65, s = 8.77	m = 38.47, s = 36.1
Vague II	m = 6.4, s = 5.31	m = 14, s = 14.28

Table 1

The following criteria was used to remove incoherent subjects: If, within a module, the subject answers ‘No’ in Question 2 and specifies a positive amount in Question 1, then the subject has committed a violation. If, within a module, the amount that the subject specifies for Question 3 is less than the amount that the subject specifies for Question 1, then the subject has committed a violation.

Table 2 is analogous to Table 1. The only difference is that every subject who committed at least two violations has been omitted.

Group	Question 1	Question 3
Clear I	m = 14.03, s = 7.75	m = 45.28, s = 36.64
Vague I	m = 5.6, s = 10.27	m = 15.76, s = 12.85
Clear II	m = 13.71, s = 9.67	m = 43.14, s = 38.33
Vague II	m = 6.4, s = 5.31	m = 14, s = 14.28

¹Clear-2 and Vague-2 were renumbered in the obvious way. That is, with Question 2 deleted, Question 3 is the second question in these modules and the numbering on the questionnaires that were used in our study reflected this change. However, for the purposes of our discussion, we will continue to use “Question 3” to denote the second question of the two modified modules.

Table 2

Table 3 is analogous to Table 1. The only difference is that every subject who committed at least one violation has been omitted.

Group	Question 1	Question 3
Clear I	m = 13.83, s = 7.88	m = 44.33, s = 36.78
Vague I	m = 6.24, s = 10.85	m = 16.79, s = 12.99
Clear II	m = 14.7, s = 11.36	m = 43.9, s = 44.92
Vague II	m = 7.33, s = 5.16	m = 15.75, s = 14.74

Table 3

Tables 2 and 3 show that the remarks following Table 1 are robust with respect to the given sensitivity analysis. In particular, our data does not support the comparative ignorance hypothesis.

In at least two of their studies, Fox and Tversky found that a comparative context increases the value of the clear(er) bet more than it decreases the value of the vague(er) bet. While these observations are not directly predicted by the comparative ignorance hypothesis, they are interesting as a supplementary hypothesis. Our data does not support the supplementary hypothesis that a comparative context increases the value of the clear(er) bet more than it decreases the value of the vague(er) bet. Table 4 summarizes some of our relevant findings:

Group	M4Q1/M1Q1	M4Q3/M1Q3
Clear I	.44	.27
Vague I	2.46	1.28
Clear II	.53	.39
Vague II	1.49	1.51

Table 4

Column M4Q1/M1Q1 shows the mean value of Question 1 in the fourth module divided by the mean value of Question 1 in the first module. Column M4Q3/M1Q3 shows the mean value of Question 3 in the fourth module divided by the mean value of Question 3 in the first module.

Recall that the questionnaires given to Clear I and Clear II have their modules arranged in terms of descending clarity, while the questionnaires given to Vague I and Vague II have their modules arranged in terms of increasing clarity. If the supplementary hypothesis were true, then one would expect – assuming that adjacent modules constitute a comparative context – that the i -th column value for Vague X would be greater than the multiplicative inverse of the i -th column value for Clear X , where $i \in \{1, 2\}$ and $X \in \{I, II\}$. This expectation is realized in only one of the four cases: in the first column the value for

Vague I is 2.46 while the multiplicative inverse of the value for Clear I is 2.27.

Table 5 is analogous to Table 4. The only difference is that every subject who committed at least two violations has been omitted.

Group	M4Q1/M1Q1	M4Q3/M1Q3
Clear I	.45	.28
Vague I	2.23	1.35
Clear II	.48	.34
Vague II	1.49	1.51

Table 5

Again, just one of four possibilities in Table 5 confirms the supplementary hypothesis: in the first column the value for Vague I is 2.23 while the multiplicative inverse of the value for Clear I is 2.22.

Table 6 is analogous to Table 4. The only difference is that every subject who committed at least one violation has been omitted.

Group	M4Q1/M1Q1	M4Q3/M1Q3
Clear I	.46	.29
Vague I	2.08	1.33
Clear II	.53	.32
Vague II	1.35	1.29

Table 6

None of four possibilities in Table 6 confirms the supplementary hypothesis.

3.1 Additional methodological issues

Our Question 1 (in the first module for all four groups) reproduces, almost unmodified, the non-comparative context utilized in Fox and Tversky's paper. The only difference is that in our study this question is the first item of a questionnaire. And there is independent research indicating that the order of presentation of items in a questionnaire can influence responses due to priming.

Fox and Weber mention the recent work by Moore [21] where he argues that the first option in a questionnaire tends to provide a context for the questions (or tasks) that follow it. So, Fox and Weber utilize these results to argue that:

We suggest that like participants in Moore's studies of decision under conflict, participants in studies of decisions under uncertainty are likely to evaluate the first prospect on a survey in a non-comparative manner, and this item is likely to provide a (comparative) context for the second (following)

items. [...] In particular, the present account predicts that when a survey contains bets that vary in their familiarity, source preference will be more pronounced for later sources than for the first source evaluated.

So, these findings seem to suggest that Question 1 in all four groups is evaluated in a non-comparative setting while subsequent questions are evaluated in a comparative setting. We assumed this in previous sections but we wanted to remark that these assumptions seem to cohere with findings in the current experimental literature.

3.2 Descriptive theories of non-expected utility under uncertainty

In 1989, in a conference in Santa Cruz, Ward Edwards posed the following two questions to a panel of distinguished decision theorists:

- 1 Do you consider SEU maximization to be the appropriate normative rule for decision -making under uncertainty?
- 2 Do you feel the experimental and observational evidence has established the fact that people do not maximize SEU; that is, that SEU maximization is not defensible as a descriptive model of the behavior of unaided decision makers?

The unanimous response was ‘yes’ to both questions. The majority of the theories of *non-expected utility* developed at least since the late 1970’s were faithful to the views revealed by the responses to Ward’s survey. In fact, most of these theories were based on the adoption of weakening s of some of the axioms of SEU plus the use of innovative techniques aimed at developing a *descriptive* theory of choice. At the same time most of these theories still presupposed that SEU was the encoding of the right normative standards for decision-making under uncertainty. One of the most influential of these descriptive theories was *prospect theory* presented in a paper by Kahneman and Tversky which appeared in *Econometrica* in 1979 [14].

Prospect theory departed from SEU in various important ways. For example, it insisted that culmination outcomes (final wealth) are not the central determinant of utility but that one should focus instead on *gains* and *loses* with respect to a reference point. In addition the theory postulated that agents evaluate uncertain prospects using *decision weights* which behave differently than the simple probability of outcomes. In particular they assumed that these decision weights are non-additive. In fact they thought

that a reliable empirical generalization is that the impact of a given event on the value of a prospect is greater when it turns an impossibility into a possibility or a possibility into a certainty than when it merely makes an uncertain event more or less probable. This is a central descriptive component of prospect theory which Fox and Tversky called in [7] *source sensitivity*. Notice that source sensitivity does not have a comparative character and it is not necessarily context-sensitive. This is an important point for our discussion in this section. We will return to it in a moment.

The first unmodified version of prospect theory has some important defects (in the view of Kahneman and Tversky themselves). For example, the appeal to non-additive decision weights is compatible with behavior which can violate first-order stochastic dominance (a prospect p' can be preferred to a prospect p , even when the cumulative probability of p of exceeding a given level of wealth is equal or greater than the probability that p' will exceed the given wealth level). And apparently this was not an empirical feature that psychologists wanted to include in their models, but rather an unintended flaw in the initial version of prospect theory. This led to the formulation of a *cumulative* version of prospect theory [26], where cumulative probabilities rather than non-cumulative probabilities are perturbed systematically by the use of decision weights. The theory was axiomatized by Wakker and Tversky [27].

Now, here we are beginning to address issues of direct interest for our study. In fact, cumulative prospect theory (CPT) is a generalization of the model of so-called ‘Choquet expected utility’ (CEU) developed by Schmeidler [24] and the so-called ‘maxmin expected utility’ (MEU) proposed by Gilboa and Schmeidler [10]. Schmeidler’s preferences can be represented by a utility function over consequences combined with a ‘Choquet capacity’. Capacities are not necessarily additive across mutually exclusive events, although they respect the condition that if event B is a subset of event A, the weight assigned to A is greater than the weight assigned to B. On the other hand, MEU represents uncertainty by utilizing a *set* of probability distributions rather than a unique distribution (or capacity, as is done in CEU). All these theories require abandoning the full force of the independence axiom of SEU.

Theories like MEU permit an explicit representation of aversion to uncertainty. This is done by adopting the minimal possible value of expected utility of a prospect when its expected utility is not completely determined. As we said above CPT is just a generalization of theories like CEU or MEU obtained by adopting the main descriptive components of PT, like

reference-points effects, etc. All these theories are potentially capable of dealing with certain versions of the Ellsberg phenomenon (MEU even includes an explicit axiom for uncertainty aversion) if not with the issue of vagueness in general (in the case of Ellsberg one has to presuppose initial conditions according to which the decision-maker has not prior stakes in the events relevant for the experiment – a condition that is satisfied in our own experimental setting).

There are in addition other descriptive theories of non-expected utility under uncertainty that can apparently be equally be applied to situations like Ellsberg's. We do not intent to be exhaustive here or to offer a review of these theories, but perhaps there are some salient accounts that can be mentioned like the theory of Machina and Schmeidler [20] or the theory proposed by Robert Nau [22] in a previous meeting of this conference. All these theories admit violations of the independent axiom of SEU and all seem to have been designed with a descriptive purpose in mind. ²

Fox and Tversky were, of course, aware of the fact that these theories exist (Tversky participated personally in the development of at least one of them) and that they are apparently applicable to study Ellsberg. But they are rather explicit about what is the range of applicability that they see in these theories. According to them (see [7], page 401) their main area of applicability is to model what they call *source sensitivity*, that is to say, the range of phenomena that is described by appealing to the non-additivity of decision weights. The main issue here is that the impact of a given event on the value of a prospect is greater when it turns an impossibility into a possibility or a possibility into a certainty, than when it merely makes an uncertain event more or less probable. In addition Tversky and Fox [25] showed that this pattern is more pronounced for uncertainty than for chance. But they also claim that source sensitivity is not the central phenomenon behind the typical pattern of preferences in the Ellsberg example. After posing the question of what are the implications of the empirical findings in [7] for individual decision making, Fox and Tversky provide their own answer to the question, which makes this point quite clearly:

To answer this question it is important to distinguish two phenomena that have

²There are other non-expected utility theories which are based on relaxing other classical axioms, like *regret theory* where transitivity is abandoned [19], [6]. But these theories are more marginally connected with Ellsberg (only some claims are made in passing by Loomes and Sudgen in their classical paper (their target is Allais) and similar considerations apply to Fishburn's paper. On the other hand, it is not clear whether the intent in these cases is either descriptive or normative.

emerged from the descriptive study of decision under uncertainty: source preference and source sensitivity. Source preference refers to the observation that choices between prospects depend not only on the degree of uncertainty but also on the source of uncertainty. Source preference is demonstrated by showing that a person prefers to bet on a proposition drawn from one source than on a proposition drawn from another source, and also prefers to bet against the first proposition than against the second. We have interpreted ambiguity aversion as a special case of source preference in which risk is preferred to uncertainty as in Ellsberg's examples. [...] The present experiments show that source preference, unlike source sensitivity, is an inherently comparative phenomenon, and it does not arise in an independent evaluation of uncertain prospects. This suggests that models based on decision weights or non-additive probabilities ([24], [27]) can accommodate source sensitivity, but that they do not provide a satisfactory account of source preference because they do not distinguish between comparative and non-comparative evaluation.

Moreover these considerations lead Fox and Tversky to claim that some theories used to model ambiguity in financial markets are probably flawed and incomplete. Incomplete, because they do not distinguish between comparative and non-comparative evaluation. And flawed because:

In particular such models are likely to overestimate the degree of ambiguity aversion in settings in which uncertain prospects are evaluated in isolation.

The claim of incompleteness might perfectly well be true. If the goal is purely descriptive it might very well be the case that comparative effects have an impact on the evaluation on ambiguous prospects. But in the light of the empirical results presented here it is less clear that the aforementioned theories are guilty of overestimating the impact of ambiguity on the evaluation of uncertain prospects in isolation. According to our results ambiguity does have a salient impact in the evaluation of vague prospects in isolation. It is an open question whether any of the descriptive theories mentioned above can be used in order to explain our data. But it is clear that our account of the 2-color problem does not need to interpret the impact of ambiguity as a special case of what Fox and Tversky call

source preference. And if any of the aforementioned theories predict that ambiguity has an impact on the evaluation of isolated prospects we will not see this deviation from SEU as a defect of the theory (at least this deviation will not necessarily be seen as an over-estimation of the impact of ambiguity in those cases).

This brings us to a more substantial issue. Almost all the theories considered above share some minimal features. On the one hand they intend to be descriptive. On the other hand almost all of them admit violations of the independence axiom. Moreover most of the authors mentioned above have not questioned the normative validity of SEU (in line with the responses to Ward's questionnaire presented at the beginning of this section). This does not mesh well neither with the line of argument defended in this paper nor with the arguments initially presented by Ellsberg which did question the normative validity of SEU.

If the normative validity of SEU is presupposed then all violations of it have to be interpreted as some sort of cognitive error, which, in turn, can be explained in terms of a corresponding psychological effect. But if one has a more open attitude towards SEU then violations of it can be classified differently according to their source. The origin of some of them might very well be traceable to the use of generally useful heuristics misfiring in unfavorable ecological contexts. But other violations can be seen in a very different light, as indications of robust behavioral patterns carrying insights about possible normative reforms of SEU. We provided some textual evidence above that Ellsberg's own views are better seen under the second point of view. It should be quite clear that the task of reforming normative standards is completely different from the task of developing a descriptive theory of choice. It is quite unfortunate that this distinction is not made more often for the purpose of classifying the so-called theories of non-expected utility. Theories of very different nature are usually lumped together under this label by the mere fact that they question some of the axioms of SEU (by adopting either weaker or stronger axioms or both). But some of the theories that abandon the axiom of ordering (mentioned in previous sections) have a clear normative purpose, while most of the theories that abandon independence have a descriptive nature (not all of them though). We are persuaded by the line of argument proposed by Ellsberg and therefore we do not see the behavioral patterns associated with the 2 and 3-color problems as an indication of some form of pathology in choice, but as a perfectly rational pattern of preferences. Moreover we see the role of experimental work in this area as having a clear bearing on the process of selecting and perfecting a theory capable of encoding norma-

tive standards of choice under uncertainty.

The task of developing descriptive theories of choice is rather different, and it is rather unclear at this point of the debate which is the canonical form that these theories should adopt. Some researchers like Rubinstein [23] or Gigerenzer [11] have questioned, for example, that these theories should be based on the adoption of particular functional forms for utility or particular functional forms representing systematic distortions in the evaluation of probability. They have proposed instead procedural accounts of choice that are more likely to have feasible computational realizations (and which are constructed on the basis of a rather different understanding of similarity than the one used in the program of heuristics and biases). We conjecture that a purely normative account of choice should be able to accommodate our data and that the kind of distortion studied in [25] can be explained as well with the help of procedural accounts of the type proposed by Rubinstein. In any case, the latter type of problem is not at the center of our inquiry, but rather the exact elucidation of the family of phenomena first discovered by Ellsberg and the normative theory(ies) capable of accommodating them.

4 Conclusions

If the comparative ignorance hypothesis is false, as our study suggests, is there a likely candidate for an alternative explanation of what we observed? We believe so, but the alternative explanation was suggested by someone who was interested in normative revisions rather than positing psychological effects that are supposed to witness human irrationality. In [4], Ellsberg offered the following informal analysis of the deviant choices he observed in connection with his own examples:

In reaching his decision, the relative weight that a conservative person will give to the question, "What is the worst expectation that might appear reasonable?" will depend on his confidence in the judgments that go into his estimated probability distribution. The less confident he is, the more he will sacrifice in terms of estimated expected pay-off to achieve a given increase in "security level"; the more confident, the greater increase in "security level" he would demand to compensate for a given drop in estimated expectation. This implies that "trades" are possible between security level and estimated expectation in his preferences, and that does seem to correspond to the observed responses. (Ellsberg, [4])

Likewise, this same analysis does seem to correspond to the observed responses in our own study. This includes trade-offs between security and expected payoffs in Question 3. Subjects are about three times more cautious in assessing these trade-offs when vagueness is present than when it is not (both in the I-groups and the II-groups).

The evidence provided by our study seems to give a stronger support to the hypothesis that agents do take into account secondary characteristics such as vagueness when assessing uncertain events in isolation. Moreover subsequent judgments seem to be highly sensible to changes in degrees of vagueness. These facts can be accommodated in a theoretical framework that takes imprecise probabilities seriously. For example, Ellsberg's own proposal can be easily modified for such a purpose. But, depending on the right articulation of a notion of security, other normative accounts (appealing to imprecise probabilities) can also be used - part of our future work intends to focus on the study of feasible theoretical options (in terms of imprecise probabilities) capable of accommodating our data. Alternative explanations in terms of psychological effects seem to lack, nevertheless, the explanatory power needed to articulate the evidence presented in this preliminary report.

References

- [1] C.C. Chow and R.K. Sarin, *Comparative ignorance and the ellsberg paradox*, Journal of Risk and Uncertainty (2001).
- [2] J. Cohen and M. Hansel, *Preferences for different combinations of chance and skill in gambling*, Nature (1959).
- [3] R.M. Dawes, G. Grankvist, and J. Leland, *Avoiding the 'ellsberg bag as avoiding a 'stacked deck possibility rather than avoiding ambiguity*, Carnegie Mellon University (2002).
- [4] D. Ellsberg, *Risk, ambiguity, and the savage axioms*, The Quarterly Journal of Economics (1961).
- [5] ———, *Risk, ambiguity and decision*, Garland Publishing, 2001.
- [6] P. Fishburn, *Utility theory for decision making*, Krieger, New York, 1979.
- [7] C.R. Fox and A. Tversky, *Ambiguity aversion and comparative ignorance*, The Quarterly Journal of Economics (1991).
- [8] C.R. Fox and M. Weber, *Ambiguity aversion, comparative ignorance, and decision context*, Organizational Behavior and Human Decision Processes (2002).
- [9] P. Gardenfors and N.E. Sahlin, *Unreliable probabilities, risk taking, and decision making*, Synthese (1982).
- [10] I. Gilboa and D. Schmeidler, *Maxmin expected utility with non-unique prior*, J. Math. Econ. (1989).
- [11] G. Gigerenzer, *Adaptive thinking: Rationality in the real world*, Oxford University Press, 2002.
- [12] C. Heath and A. Tversky, *Preference and belief: Ambiguity and competence in choice under uncertainty*, Journal of Risk and Uncertainty (1991).
- [13] W.C. Howell, *Uncertainty from internal and external sources: A clear case of overconfidence*, Journal of Experimental Psychology (1971).
- [14] D. Kahneman and A. Tversky, *Prospect theory: an analysis of decision under risk*, Econometrica (1979).
- [15] J.M. Keynes, *A treatise on probability*, MacMillan, 1921.
- [16] F.H. Knight, *Risk, uncertainty and profit*, Houghton-Mifflin, 1921.
- [17] I. Levi, *On indeterminate probabilities*, Journal of Philosophy (1974).
- [18] ———, *The paradoxes of allais and ellsberg*, Economics and Philosophy (1986).
- [19] G. Loomes and R. Sudgen, *Regret theory: an alternative theory of rational choice under uncertainty*, The Economic Journal (1982).
- [20] M. Machina and D. Schmeidler, *A more robust definition of subjective probability*, Econometrica (1992).
- [21] D.A. Moore, *Order effects in preference judgments: Evidence for context dependence in the generation of preferences*, Organization Behavior and Human Decision Processes (1999).
- [22] R. Nau, *Uncertainty aversion with second-order utilities and probabilities*, 2nd International Symposium on Imprecise Probabilities and Their Applications (<http://www.sipta.org/isipta01/proceedings/063.html>) (2001).

- [23] A. Rubinstein, *Similarity and decision making under risk (is there a utility theory resolution of the allais paradox?)*, Journal of Economic Theory (1988).
- [24] D. Schmeidler, *Subjective probability and expected utility without additivity*, Econometrica (1989).
- [25] A. Tversky and C. Fox, *Weighing risk and uncertainty*, Psychological Review (1995).
- [26] A. Tversky and D. Kahneman, *Advances in prospect theory: Cumulative representation of uncertainty*, Journal of Risk and Uncertainty (1992).
- [27] P. Wakker and A. Tversky, *An axiomatization of cumulative prospect theory*, Journal of Risk and Uncertainty (1993).