

On the Explanatory Power of Indeterminate Probabilities

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

Jeffrey Helzner
Columbia University
jh2239@columbia.edu

Abstract

Building on work that we reported in [1] we revisit the claims made by Fox and Tversky in [3] concerning their *comparative ignorance* hypothesis for decision making under uncertainty.

Keywords. normative, descriptive, indeterminate, rationality.

1 Introduction

The purpose of this paper is to report on recent developments in the research program that we introduced in [1] and continued in [2]. The motivating questions behind this program concern the extent to which certain normative theories of decision making that are based on indeterminate probabilities can be used to rationalize some observed deviations from the orthodox Bayesian account. This program may be contrasted with the so-called ‘heuristics and biases’ program which seeks descriptive theories of decision making that are capable of accommodating such deviations. Where our program is focused on the explanatory power of indeterminate probabilities as they are employed in certain normative theories, the heuristics and biases program introduces clearly non-normative explanatory devices such as psychological effects in order to accommodate certain deviations from orthodox Bayesian account.

Fox and Tversky’s ‘comparative ignorance’ hypothesis of [3] is an important example of the sort of theoretical work that has been advanced within the heuristics and biases framework. Roughly, the basic idea behind comparative ignorance is that ‘uncertainty aversion’ is mainly driven by comparative contexts in which the decision maker is made aware of their lack of knowledge concerning a given uncertain event as a result of the salience of another uncertain event about which they are better informed. Fox and Tversky investigate the comparative ignorance hypothesis through a series

of experiments that employ both ‘clear’ and ‘vague’ prospects in isolation and jointly. Following Fox and Tversky’s terminology we will say that a comparative context obtains when the subject is presented with both a clear and a vague prospect.

The following illustrates a type of prospect that Fox and Tversky use in their work: An urn has been filled with 100 balls, where m of the balls are known to be solid black and $n \leq 100 - m$ are known to be solid white. What is the most that you would be willing to pay for a ticket that pays \$100 if a given random selection from the urn yields a black ball and pays \$0 dollars if the random selection yields a white ball? Prospects of this type for which $m + n = 100$ are said to be clear while those for which $m + n < 100$ are said to be vague.

Now suppose that we fix two prospects of the indicated type. Prospect A is clear and is determined by setting $m = 50 = n$. Prospect B is vague and is determined by setting $m = 0 = n$. According to the comparative ignorance hypothesis subjects who are presented with both A and B , thereby constituting a comparative context, will tend to exhibit a significant difference in their willingness to pay for these prospects while the difference between the maximum purchase price for those subjects who are offered A in isolation and those who are offered B in isolation will tend to be relatively insignificant.

In [1] we employed a rather novel methodology whereby subjects were asked to price their ticket on mixtures of the basic chance setups that were considered in Fox and Tversky’s original experiment. For example, subjects were asked to state the most they would be willing to pay for their ticket given that the payoff would be determined by a two-stage chance setup where the first stage consists of a flip of a fair coin and the second stage, which depends on the outcome of the first, consists of a random draw from either the 50:50 urn or the urn that has a completely unknown ratio of black balls to white balls. With this

methodology we introduced gradations between Fox and Tversky’s clear and vague bets; increasing the bias of the coin in the first stage in favor of the 50:50 urn results in a more clear bet while increasing the bias of the coin in the first stage in favor of the completely indeterminate urn results in a more vague bet. As we reported in [1] increasing amounts of vagueness had significance for the maximum buying prices even in the absence of a comparative context of the sort that was discussed by Fox and Tversky.

2 Experiment on Mixtures of Chance Setups

In our initial use of mixtures of chance setups we assumed that these setups were, at least in principle, reducible to one-stage setups. Given that an overarching goal of our project has been to investigate the explanatory power of indeterminate probabilities we have a clear interest in whether or not a play on a mixed chance setup can be exchanged for a play on a particular urn for which the subject is given some, although not necessarily complete, information concerning the ratio of black balls to white balls. We now report on study that takes some first steps towards an understanding of this reduction.

Our subjects, 56 undergraduates at Carnegie Mellon University, were presented with a questionnaire that began with the following description of the relevant chance setups:

Urn A contains exactly 100 balls. 50 of these balls are solid black and the remaining 50 are solid white.

Urn B contains exactly 100 balls. Each of these balls is either solid black or solid white, although the ratio of black balls to white balls is unknown.

Urn X contains exactly 100 balls. Each of these balls is either solid black or solid white. Further assumptions concerning this urn will be considered in the questions below.

The next item on the questionnaire was the following presentation of the alternatives that the subjects were asked to consider:

Alternative 1 We flip a fair coin. If the coin lands heads, then we draw a ball at random from Urn A. If the coin lands tails, then we draw a ball at random from Urn B. In either case, if the ball that is drawn is black, then you get \$100. However, if the ball that is drawn is white, then you get \$0.

Alternative 2 We draw a ball at random from Urn X. If the ball that is drawn is black, then you get \$100. If the ball that is drawn is white, then you get \$0.

Finally, we attempted to elicit a reduction with the following two questions:

Question 1: What is the smallest number m , between 0 and 100, such that if you were to learn that X contains at least m black balls then you would be willing to choose Alternative 2 when offered a choice between Alternative 1 and Alternative 2?

Question 2: Let m be your answer from Question 1. Assume that all you know about the distribution of black balls and white balls in Urn X is that there are at least m black balls in Urn X so that, according to your answer to Question 1, you would be willing to choose Alternative 2 when offered a choice between Alternative 1 and Alternative 2. Is there a number n , between 0 and $100 - m$, such that if you were to learn that Urn X contains at least n white balls then you would no longer be willing to choose Alternative 2 when offered a choice between Alternative 1 and Alternative 2? If there is such a number, then write the least such number in the space below.

There are at least two rather obvious theoretical candidates that might be considered in connection with the questionnaire that has just been presented. According to the first of these, a reduction is given by the following operation:

$$U \oplus_{\lambda} V = \{\lambda p + (1 - \lambda)q \mid p \in U \text{ and } q \in V\} \quad (1)$$

where U and V are each sets of probability distributions over a common state space and $\lambda \in [0, 1]$ is a mixture weight. Interpret A as the set that contains exactly one distribution, namely the one that assigns a probability to drawing a black ball that is equal to that of drawing a white ball. Interpret B as the set of all distributions on {Black, White}. Taking λ as $\frac{1}{2}$, $A \oplus_{\frac{1}{2}} B$ evaluates to the set of all distributions p on the indicated set of states such that $p(\text{Black}) \geq \frac{1}{4}$ and $p(\text{White}) \geq \frac{1}{4}$. This suggests a reduction of the two-stage chance setup to a single-stage setup where the subject is told that there are at least 25 black balls and at least 25 white balls in the urn. For our second theoretical candidate imagine a subject who applies the principle of insufficient reason when considering B , the maximally indeterminate urn, and thus interprets the flip of the coin as leading to a play on a 50:50 urn in either case. This second account suggests a reduction of the two-stage setup to a single stage setup

	> 37.5	< 37.5
Black min	39	17
White max	47	9

Table 1:

> 37.5	< 37.5
34	2

Table 2:

where the subject is told that the urn has 50 black balls and 50 white balls.

It turns out that relatively few of the subjects were in complete agreement with either of the theoretical candidates. As noted, the first (second) theoretical candidate suggests a value of 25 (50) as the minimum number of known black balls required for Alternative 2 to be admissible and the first (second) theoretical candidate suggests a value of 25 (50) as the maximum number of known white balls such that Alternative 2 retains its status as an admissible option.¹ Now, if we suppose that each of these candidates is playing a role to some degree, then we can try to consider the extent to which one of the two seems to dominate by making a cut at 37.5, i.e. the midpoint ($\frac{25+50}{2}$) between the theoretical predictions regarding each of the two bounds.

Table 1 shows the number of subjects who reported a value above (below) the midpoint for each of the two bounds. Table 1 suggests that the second model, the one based on the principle of insufficient reason, is dominant, at least when the two questions are considered individually. Table 2 shows the breakdown when fit is considered with respect to both of the bounds. The first column of Table 2 shows the number of subjects who gave values above 37.5 for both of the bounds (i.e. the minimum for black and the maximum for white). Similarly, the second column shows the number of subjects who gave values below 37.5 for both of the bounds. Again the model that is based on the principle of insufficient reason appears dominant, 34 of 56 compared to 2 of 56.

While the analysis above suggests that the data favors the account based on the principle of indifference, it is important to note that relatively few subjects returned values that are in complete agreement with this account. As a possible explanation of this

¹This first set of values (i.e. the minimum number of known black balls) may be computed directly from Question 1, while the second set of values (i.e. the maximum number of known white balls) can be computed from Question 2 as $n - 1$ if a value is supplied and $100 - m$ if no value is supplied by the subject.

some might suggest that the mixed chance setup is itself a comparative context in that it makes salient a comparison between the maximally indeterminate urn and the 50:50 urn. To be sure, this is not quite a comparative context in the sense of Fox and Tversky: there is but one alternative being played against the mixed chance setup. Moreover, it is unclear how Fox and Tversky can in general interpret the upper and lower bounds that are reported. Of course we are open to interpreting these bounds as upper and lower probabilities for a potential credal state, but such an interpretation does not seem to be an option for Fox and Tversky. Nonetheless, let us allow a very generous interpretation of ‘comparative context’ so that the results of [1], which employed mixed chance setups, can be countered by objecting that we did in fact employ a comparative context. The following studies attempt to reinforce our point without appealing to mixed chance setups.

3 Two Experiments That Do Not Use Mixtures of Chance Setups

Fox and Tversky predict essentially Bayesian behavior in the absence of a comparative context. Moreover, in keeping with much of the heuristics and biases tradition, they seem to interpret all deviations from the Bayesian standard as instances of irrationality; deviations from this standard are predicted when the subject is under the spell of various psychological effects, which in the sort of cases that we have been considering are those resulting from the presence of a comparative context. While our primary focus in [1] was to argue against some of the core claims in [3], we will now consider data that has direct relevance to the manner in which Fox and Tversky seem to interpret deviations from the Bayesian standard.

Using a protocol that was derived from an example in [4] we attempted to ascertain the extent to which violations of the Bayesian standard could be accommodated by certain normative alternatives based on indeterminate probabilities. For the purposes of these studies we focused on the following decision rules:

E-admissibility: Assume that the decision maker’s credal state can be represented by a set P of probability distributions. If X is a set of alternatives, then a is *E-admissible* in X iff $a \in X$ and there is some distribution in $p \in P$ for which $E_p[a] \geq E_p[b]$ for all $b \in X$, where $E_p[x]$ is the expected utility of x against distribution p .

MMEU: Assume that the decision maker’s credal state can be represented by a set P of probability distributions. For each alternative x , let x_- be

the greatest lower bound of $\{E_p[x] \mid p \in P\}$. If X is a set of alternatives, then $x \in X$ satisfies the *maximin criterion for expected utility* (MMEU) on X iff $x_- \geq y_-$ for all $y \in X$.

E-admissibility followed by MMEU: Assume that the decision makers credal state can be represented by a set P of probability distributions. If X is a set of alternatives, then a is admissible in X according to this rule iff a satisfies MMEU on the set of alternatives that are E-admissible in X .

E-admissibility is discussed in [6] where it is taken as the first tier in Levi’s two-tiered decision theory. Alternatively, taken as a free standing decision rule, E-admissibility corresponds to Levi’s theory when the security, the second tier, is vacuous. MMEU has a long history in the statistics literature, e.g. in discussions of ‘gamma minimax’, and continues to receive attention in decision contexts [4, 5]. The third criterion is essentially an instance of Levi’s decision theory where the second tier security rule is given by MMEU.

3.1 Study 1

Our subjects, 56 undergraduates at Carnegie Mellon University, were presented with a questionnaire that began with the following description of the underlying chance setup.

An urn has been filled with several balls, each of which is either solid black or solid white. While the exact ratio of black balls to white balls is unknown, the following statistical information is available:

Black The probability of selecting a black ball on a single random draw from the urn is at least %40 but not more than %60.

White The probability of selecting a white ball on a single random draw from the urn is at least %40 but not more than %60.

The next section of the questionnaire introduced the following choice problem, which, as noted above, is based on an example from [4]:

Consider the three alternatives in the table below. Note that the payoffs for these alternatives are in dollars. So, for example, Alternative pays \$-10 if a black ball is drawn (i.e. you lose \$10 if a black ball is drawn from the urn) and pays \$12 dollars if a white ball is drawn.

	Black	White
A	-10	12
B	11	-9
C	0	0

A,B	B	C	Other
12	9	22	13

Table 3:

In the final section of the questionnaire the subjects were given the following prompt and then asked to indicate the alternatives that they would be willing to choose.

Suppose that you are offered the opportunity to specify the alternatives above that you are willing to choose, with the understanding that we will pick one of these alternatives *before the random selection from the urn* and you will receive the winnings, or pay the losses, that are generated by the alternative that we pick.

Before turning to the data from this initial study, let us apply the three decision rules from the previous section to the choice problem that is presented in the above questionnaire. Assume that utilities are determinate and linear in dollars. Assume that the agent’s credal state is given by the description at the beginning of the protocol. That is, assume that the agent’s credal state can be represented by the set $P = \{p : .4 \leq p(\text{Black}) \leq .6\}$. Under these assumptions it follows that A and B are the only E-admissible alternatives in $\{A, B, C\}$ while C is the only alternative that satisfies MMEU in $\{A, B, C\}$. The two-tiered rule, E-admissibility followed by MMEU, counts B as the lone admissible alternative since B is the only alternative that satisfies MMEU in $\{A, B\}$.

Table 3 shows that those subjects who regarded C as uniquely admissible constitute the largest group by a rather wide margin, with the total for C being approximately equal to the combined totals for A, B and B . It is worth noting that C is the only alternative that fails to be a bayes solution under the assumption that utilities in this range are determinate and essentially linear in dollars. Given that this is a noncomparative problem, C ’s dominant position seems to disconfirm Fox and Tversky’s prediction of essentially Bayesian behavior in absence of a comparative context. This point can be strengthened if we note that Fox and Tversky seem to have in mind that the appropriate Bayesian model for predicting behavior in noncomparative contexts is one that appeals to the principle of insufficient reason. That is, Fox and Tversky seem to predict that subjects who are faced with the given noncomparative choice problem will choose as though they are maximizing expectations against the distribution that assigns the two states an equal probability. Assuming that utilities are determinate and linear in dollars, subjects who are in accordance with this prediction must be willing to choose A and B . Hence,

in addition to those who judged C to be uniquely admissible, those who judged B to be uniquely admissible fail to confirm Fox and Tversky's prediction for noncomparative choice under uncertainty.

Although their presence seems to disconfirm Fox and Tversky's predictions, at least under the assumption that utilities are determinate and linear in dollars, those who judged either B or C to be uniquely admissible in the triple are consistent with one of the normative alternatives discussed above under this very same linearity assumption. Of course we recognize that by appealing to other considerations, e.g. non-linear utilities or perhaps various psychological effects, one might formulate decidedly non-normative decision models that are capable of reproducing the admissibility judgments that were reported by these subjects. It is for this reason that we decided to conduct a further investigation in order to determine the extent to which subjects reasoned in the manner suggested by the normative theory that reproduced their admissible choices. Details of this second study are the subject of the next section.

3.2 Study 2

Our subjects, 27 undergraduates at Carnegie Mellon University, were presented with a questionnaire that began exactly as the one that was employed in the previous study but continued with the following illustrations of each of the three decision rules that were discussed at the beginning of Section 3.

Albert's reasoning: Note that the only probability distributions that are consistent with the information that is given are those for which the probability assigned to drawing a black (white) chip is at least .4 and no more than .6. Among the distributions that satisfy these conditions, there are some for which A maximizes expected value. For example, if the probability of drawing a black ball is .4, and so the probability of drawing a white ball is .6, the following table gives the expected value of each of the alternatives.

	$p(\text{Black}) = .4$	$p(\text{White}) = .6$	Expected Val.
A	-10	12	3.2
B	11	-9	-1.0
C	0	0	0

From the table it is clear that A maximizes expected value against the probability distribution that assigns a probability of .6 to drawing a white ball. Similarly, B maximizes expected value against the probability distribution that assigns a probability of .6 to drawing a black ball and a probability of .4 to drawing a white ball. On the other hand there is no distribution that is consistent with the information that is given and against which C maximizes expected value. With this reasoning I elimi-

nated C from further consideration. I would be willing to choose A and B, but not C.

Bob's reasoning: Well, I eliminated C along the same lines as Albert suggested, but then I appealed to some additional considerations. Since the minimal expected value of A is -1.2 , which occurs when the probability assigned to drawing a black ball is .6, and the minimal expected value of B is -1.0 , which occurs when the probability assigned to drawing a white ball is .6, I decided to eliminate A from further consideration. I would be willing to choose B, but not A or C.

Carol's reasoning: My reasoning was essentially like Bob's, except for the part where he followed Albert. That is, I simply considered the minimal expectation of each of the three alternatives. Since the minimal expected value of A is -1.2 and the minimal expected value of B is -1.0 , while the minimal expected value of C is 0, I eliminated A and B from further consideration. C has the largest minimal expectation. So, I would be willing to choose C, but not A or B.

Finally, the subjects in this study were asked to indicate their level of agreement with each of the statements below. We instructed the subjects to indicate their level of agreement on a scale from 1 to 10 (i.e. 1, 2, 3, ..., 10) with 1 being 'not at all' and 10 being 'as much as possible':

- Albert's reasoning is compelling.
- Bob's reasoning is compelling.
- Carol's reasoning is compelling.
- Albert's reasoning resembles the reasoning that I used in formulating my own response to the question.
- Bob's reasoning resembles the reasoning that I used in formulating my own response to the question.
- Carol's reasoning resembles the reasoning that I used in formulating my own response to the question.

First, before turning the data obtained from the additional questions, we recall that the subjects in this second study also answered the questions that were given to those in the first study. Table 4 shows the breakdown of this group of subjects in terms of the same partition that was employed in Table 3. As was reported in Table 3 in connection with the first study, Table 4 shows that the group of subjects who judged C to be uniquely admissible in the triple is the largest of the four groups and, as before, is roughly the size of the groups for A, B and B combined.

A,B	B	C	Other
7	3	12	5

Table 4:

It is important to remember that the first part of the questionnaire that was employed in Study 2 is identical to the questionnaire that was used in Study 1. One of the referees who commented on an earlier version of this paper suggested that the additional questions that were used in Study 2 might have led the subjects. There are at least two reasons to believe that this is not the case. First, the additional questions, i.e those concerning the three character sketches, were posed at the end of the questionnaire. The subjects were instructed to respond to the questions in the order that they were presented and were told not to go back to revise their answers to earlier questions. Second, Table 4, which shows the data from the part of the questionnaire that duplicated what was used in Study 1, suggests a very similar breakdown to what was observed in Study 1.

Returning to the matter that prompted this second study, we note that 15 of the 27 subjects reported *B* or *C*. The issue that prompted this second study concerns the extent to which these subjects determine admissibility by appealing to the considerations that are suggested by one of three non-Bayesian, normative rules described above. In terms of the additional questions that were employed in this second study we can attempt to address this question by isolating those subjects who reported a high level of resemblance between their own reasoning and the appropriate non-Bayesian norm. Interpreting a high level of resemblance to be a value of 8 or above for the subject's response to the relevant question we observed that 11 of these 15 appealed to considerations that had a high level of resemblance to those suggested by the appropriate non-Bayesian norm. Finally, although the numbers are getting rather small at this point it is perhaps worth noting that although *A, B* is consistent with Fox and Tversky's predictions the majority of the subjects in the *A, B* group reported that their reasoning had a high resemblance to Albert's *E*-admissibility considerations.

4 Conditional Support

Subjects evaluate the resemblance of their own form of reasoning to the theories exemplified by Albert, Bob and Carol's reasoning at the end of the questionnaire. Previous to this, and after receiving the information about Albert, Bob and Carol, they assess the validity of these theories. One minimal desideratum here is

that subjects who judge their own reasoning to have a strong resemblance to theory X (with X varying over the theories advanced by Albert, Bob and Carol) rank the theory X with a score superior to at least 5 in the scale from 0 to 10. Otherwise we would have a situation where subjects see themselves as judging according to a theory that they themselves consider to have dubious validity.

Not all subjects obey these minimal desiderata and we propose to filter them out in order to consider unconditional and conditional support for the three theories under consideration. In particular there is a subject who chooses C and sees himself as choosing according to Carol's considerations but gives Carol's theory a score of 4.

If we consider unconditional support for C after this subject is eliminated from the pool of respondents the average unconditional support for C has a value of 8.54 (in comparison with a value of 8.1 before eliminating subjects who do not obey the aforementioned desideratum).

There is also a separate measure of interest which is given by the amount of support that a theory X received conditional on the fact that the subject chooses what X recommends and that the subject sees himself as choosing in accordance to X. We will consider that a subject sees herself as choosing in accordance with X if she ranks X as resembling her form of reasoning with a score of at least 8.

The average conditional support for C in these circumstances has a value of 9.25. Similarly the average conditional support for A, B has a value of 8.25. And the corresponding average conditional (and unconditional) support for B has a maximal value of 10. So, all the values of average conditional support are relatively high.

The average unconditional support for A, B is, nevertheless, lower than the average conditional support (7.71) indicating that there are some subjects who choose A, B but do not see themselves as choosing in accordance with Albert's form of reasoning. It would be interesting in future research to consider alternative forms of reasoning consistent with choosing A, B even when they might not be articulated in terms of indeterminate probabilities (applications of the principle of insufficient reasoning might be a salient option here).

In the case of option C the gap between unconditional and conditional support is less significant (8.54 after sensitivity analysis as opposed to 9.25) indicating that this form of non-Bayesian reasoning is very robust. Finally there is no gap between average unconditional

and conditional support in the case of B. This form of non-Bayesian reasoning occurs in a minority of cases as opposed to A, B and B, but it is supported in a very strong manner when it occurs.

5 Future work

The comparative ignorance hypothesis advanced by Fox and Tversky explains deviations from Bayesian behavior in cases where there is indeterminacy in terms of a psychological effect, namely *uncertainty aversion* driven by comparative contexts. But as we tried to make clear in this paper there are frequent cases of decision contexts where there is indeterminacy but no comparison is being made. The scenario in Study 1 is such a case. As we stressed above this is a case where there is no comparison between clear and vague bets. All bets are vague. In a situation of this sort it seems that there is no psychological effect, at least along the lines that were suggested by Fox and Tversky's account, that one might invoke to predict a deviation from Bayesian standards of rationality.

There are, nevertheless, several decision models that take indeterminacy seriously and might be used to explain the behavior verified in the experiments that we have reported. Of course there could be other models that take indeterminacy into account but in a way that is very different from what is suggested in the decision rules that we have considered. Alternatively, there might be an entirely different psychological effect that is driving the behavior of subjects. The existence of all these possibilities is what motivated our second experiment, where some theoretical options were presented to the subjects for their appraisal. Subjects had the option of saying that none of the presented options represented their reasoning adequately. Nevertheless we verified that 11 out of 15 subjects selected one of the theoretical options as closely resembling their reasoning. So, this seems to indicate that one of the theoretical options that takes indeterminacy seriously (MMEU) figures among the reasoning strategies of actual subjects.

Is it possible that psychological effects that have nothing to do with indeterminacy motivate the non-Bayesian behavior verified in the experiments? One referee pointed out that the fact that our example uses negative payoffs might be the cause of some of the behavior observed in the experiments. The idea is that agents might be motivated by loss aversion and that this explains the selection of option C in Study 1 (2). This nevertheless does not explain why subjects selected Carol's reasoning as resembling as much as possible their own reasoning. One needs to assume here that a majority of subjects were mistaken in as-

sessing their own reasoning.

One experiment that can settle the issue (as suggested by the referee) is to run a version of Study 1 (2) where 15 dollars is added uniformly to all payoffs in the matrix used in both experiments. The referee predicts that in this case option C will lose its appeal and that options A and B (a Bayes solution) would be chosen. Notice that even if this behavior were observed this solution is also compatible with Albert's reasoning (E-admissibility). To settle this issue we propose to add a theoretical option along the lines of the principle of insufficient reason to the salient theoretical options offered to the subjects in a new version of experiment two. Fox and Tversky seemed to have predicted that in a situation of this sort agents will appeal to insufficient reasoning. Other methods of dealing with indeterminacy (like Albert's reasoning) remain possible as well. So, the problem of determining which one of these methods constitutes an empirically robust response to indeterminacy remains as open in this new experimental set up as it was in the scenario investigated in this paper.

6 Conclusions

In Section 2 we reported on an experiment that was conducted in order to investigate the manner in which subjects reduce mixtures of chance setups, of the sort that we employed in [1], to indeterminate probabilities. This was important to us because part of our overall research program is an exploration of the explanatory power of indeterminate probabilities, especially as this stands in contrast to the purely descriptive agenda that is articulated within the heuristics and biases paradigm. As we discussed in Section 2, the data that were generated by this experiment raised the possibility that mixtures of chance setups might constitute a comparative context of sorts. If so, then such a thing could be offered as an objection to the arguments that we advanced in [1], e.g. one could object that we had smuggled in a comparative context by using mixtures of chance setups. Anticipating this objection we conducted the two experiments that are reported in Section 3. These two experiments address the core of Fox and Tversky's claims without appealing to mixtures of chance setups. The results from the first of these studies suggests a significant amount of non-Bayesian behavior occurring in a noncomparative context. We were able to rationalize much of this non-Bayesian behavior in terms of three well-known normative rules that are based on indeterminate probabilities. The second study in Section 3 suggests that such rationalizations of the indicated non-Bayesian behavior are not merely of the 'as if' variety but rather approximate a sub-

stantial portion of the reasoning that is driving this behavior. Thus, despite the claims of Fox and Tversky, it appears that there is a significant amount of non-Bayesian behavior even in the absence of a comparative context.

References

- [1] H. Arlo Costa and J. Helzner, *Comparative ignorance and the ellsberg phenomenon*, Proceedings of the Fourth International Symposium on Imprecise Probabilities and Their Applications (F. Cozman, R. Nau, and T. Seidenfeld, eds.), 2005.
- [2] ———, *Risk, indeterminacy and decision*, Proceedings of the ESSLI 2006 Workshop on Rationality and Knowledge (R. Parikh and S. Artemov, eds.), 2006.
- [3] C.R. Fox and A. Tversky, *Ambiguity aversion and comparative ignorance*, The Quarterly Journal of Economics (1991).
- [4] P. Gardenfors and N.E. Sahlin, *Unreliable probabilities, risk taking, and decision making*, Synthese (1982).
- [5] I. Ghilboa and D. Schmeidler, *Maxmin expected utility with non-unique prior*, J. Math. Econ. (1989).
- [6] I. Levi, *On indeterminate probabilities*, Journal of Philosophy (1974).