

RECOGNIZING AMBIGUITY HOW LACK OF INFORMATION SCARES US

Mark Clements
Columbia University

I. Abstract

In this paper, I will examine two different approaches to an experimental decision problem posed by Craig Fox and Amos Tversky. This decision problem is intended to illustrate the phenomenon called “ambiguity aversion” that is observed in decision situations under uncertainty. These situations occur when an agent is faced with a decision problem where the probabilities are not specified in advance or readily assessed based on the given information in the decision problem. Because ambiguity aversion in decisions under uncertainty is one of the most fundamental problems of traditional decision theory¹, in attempting to give an account of rational decision making Fox and Tversky offer a hypothesis to account for ambiguity aversion and explain the decision situations that exhibit this problem. After examining their explanation for ambiguity aversion, I will consider another possible approach to these decision problems under uncertainty viz. the recognition heuristic, which has

1 By “traditional decision theory” I mean subjective expected utility maximization.

been offered as a decision tool by Daniel Goldstein and Gerd Gigerenzer under an alternate approach to decision theory called “bounded rationality.”² The purpose of this exercise is to examine how each of these approaches accounts for decision behavior when faced with the same decision problem. In doing so, we will be able to determine how relevant these competing theories are to each other and discover the limitations of the heuristic approach in pursuing a comprehensive model for rational decision making.

II. Framework and Hypothesis

Before proceeding, we need to establish the framework for the hypothesis that is under question here. Fox and Tversky are attempting to identify decision situations where agents exhibit ambiguity aversion. They claim that comparative ignorance effects cause ambiguity aversion and that this is an instance of preference reversal.³ Fox and Tversky don’t attempt to solve the problem of preference reversal; rather, they pose a hypothesis to describe the decision situations where ambiguity aversion manifests. They successfully support their hypothesis in experimental decision situations by restricting ambiguity aversion to situations where comparative ignorance exists. However, this only identifies the

2 Simon, 1955

3 It’s debatable if preference reversal is confined to comparative ignorance and ambiguity aversion, but this seems to be the approach taken by Fox and Tversky.

domain of relevance for ambiguity aversion while ignoring the more general problem of preference reversal in traditional decision theory.

This reversal of preferences is one of the major problems faced by traditional expected utility theory. It is observed in the Ellsberg problem, where agents systemically violate one of the basic axioms of subjective expected utility theory: Savage's "sure thing principle," or the "independence" axiom. I'm not going to digress into an explanation of this problem because it is a prominent issue in the literature, so prior familiarity with the issues posed by Ellsberg will be assumed here. Essentially, Fox and Tversky claim that the preference reversal is attributed to ambiguity aversion under comparative ignorance.

Because Fox and Tversky fail to address the broader difficulties posed by ambiguity aversion, we might wonder if the recognition heuristic of bounded rationality proposed by Goldstein and Gigerenzer could resolve this problem. If so, how would it resolve the problem? If the heuristic model proposed by Goldstein and Gigerenzer can be applied to Fox and Tversky's decision experiment and develop the same empirical trends⁴, then decision patterns that exhibit preference reversal can be accounted for in a theory of rational decision making. The recognition heuristic would provide an alternative method,

4 That is, the heuristic will also instruct agents to choose the clear bet, which is the decision exhibited in Fox and Tversky's experiment.

which would be devoid of the problematic axioms and utility maximization demanded in the traditional theory for explaining such decision problems under uncertainty. I will attempt to apply the recognition heuristic to Fox and Tversky's experimental decision problem to see if it is able to give the same decision outcomes. If this application is successful in this case, then there might be good reason for using this heuristic to supplement the traditional decision theory in situations under uncertainty that illicit irrational choice behavior due to axiom violation. If this attempt is not successful, then the relevance and usefulness of such heuristics in a comprehensive model for decision making must be examined.

III. Fox and Tversky – Ambiguity Aversion

I will begin my analysis by examining Fox and Tversky's experimental setup and the decision problem that is being observed. As stated above, they propose a hypothesis to account for decision situations that exhibit ambiguity aversion. This is their "comparative ignorance hypothesis", which states "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."⁵ This is the hypothesis they are testing in their experiment, which presents people with the following decision

5 Fox and Tversky, 588

problem:

Imagine that there is a bag on the table (Bag A) filled with exactly 50 red poker chips and 50 black poker chips, and a second bag (Bag B) filled with 100 poker chips that are red and black, but you do not know their relative proportion. Suppose that you are offered a ticket to a game that is to be played as follows: First, you are to guess a color (red or black). Next, without looking, you are to draw a poker chip out of one of the bags. If the color that you draw is the same as the one you predicted, then you will win \$100; otherwise you win nothing. What is the most you would pay for a ticket to play such a game for each of the bags? (\$0-\$100)

<u>Bag A</u>	<u>Bag B</u>
50 Red Chips	? Red Chips
50 Black Chips	? Black Chips
100 Total Chips	100 Total Chips

The most I would be willing to pay for a ticket to Bag A (50 red; 50 black) is:___

The most I would be willing to pay for a ticket to Bag B (? red; ? black) is:_____6

Fox and Tversky conducted this experiment with three groups of people. The first group evaluated the clear and vague bets in a comparative situation as shown above. The other two groups evaluated the bets in a non-comparative

situation with one evaluating the clear bet alone, and the other, the vague bet alone. The results of their experiment show that in the comparative group, people priced the clear bet significantly higher than the vague bet; in the non-comparative groups, there wasn't a significant price difference between these two bets.⁷ The decision pattern exhibited in this experiment supports their hypothesis that in decision situations of comparative ignorance, ambiguity aversion is exhibited by people preferring the clear bet to the vague bet.⁸

In verifying their hypothesis, Fox and Tversky conclude that comparative situations between one clear option and one vague option exhibit ambiguity aversion, and that this accounts for preference reversal in cases such as Ellsberg's problem.⁹ They further conclude that traditional decision theory "requires that the comparative and non-comparative evaluations will coincide" (by virtue of the independence axiom), but it does not "provide a method for reconciling inconsistent preferences."¹⁰ By successfully identifying choice situations of comparative ignorance as exhibiting ambiguity aversion, their results are beneficial to traditional decision theory because they restrict the domain of relevance of the preference

7 *Ibid.*, 589

8 Preference in the sense that they were willing to wager more on the clear bet over the vague bet. They go on to conduct several other experiments that show similar results in support of this hypothesis.

9 *Ibid.*, 600

10 *Ibid.*

6 *Ibid.*

reversal problem to decision situations of this type. Although this identification is valuable for traditional decision theory, it doesn't solve the fundamental problem illustrated by Ellsberg. Could bounded rationality's alternative approach to decision making provide a solution to this problem? Possibly, but we would have to determine which tool in the "adaptive toolbox" is appropriate to use in this situation. Now that the decision problem of Fox and Tversky has been identified I will define the recognition heuristic as described by Goldstein and Gigerenzer, provide a justification for attempting to apply it here, and then see if it is, in fact, applicable in this case.

IV. The Recognition Heuristic

In the book Simple Heuristics that Make Us Smart, Daniel Goldstein and Gerd Gigerenzer outline a simple heuristic that describes decision patterns under a specific domain of relevance. To see if we can apply this heuristic to the decision problem presented by Fox and Tversky, it is important to understand how the heuristic works and determine the necessary conditions for implementation.

Since the heuristic functions based on the concept of 'recognition,' this is the first concept that they define. They claim that the term "recognition" has been used in many contexts, but for the purposes of this heuristic, the term is based on the simple binary relationship between the

novel and the previously experienced.¹¹ Simply put, "recognized objects" are objects we've experienced before and "unrecognized objects" are novel objects. With this understanding of 'recognition,' Goldstein and Gigerenzer explore the mechanics of the heuristic by asking an individual to consider the decision problem of "inferring which of two objects has a higher value on some criterion."¹² In evaluating this decision problem, the recognition heuristic says that "if one of two objects is recognized and the other is not, then infer that the recognized object has the higher value."¹³

Goldstein and Gigerenzer demonstrate the use of their heuristic in an experiment involving German and American students deciding which particular American city has the higher population. In their example, the two objects are San Diego and San Antonio, and the criterion is population. The question is then posed: which city, San Diego or San Antonio, has the higher population? According to their results, the German students were as accurate, if not more so, than the American students at deciding which of the two American cities had the larger population. These results are interesting because the German students were more successful in reaching a correct conclusion. The idea is that they were able to use the recognition heuristic

11 Goldstein and Gigerenzer, 38

12 *Ibid.*, 41

13 *Ibid.*

more often than the American students because many of them recognized San Diego, but not San Antonio. From this, the authors describe what they call the “less-is-more effect” which states that more information doesn’t necessarily create accuracy and, in fact, there is a certain domain of “ignorance” over which the heuristic is most effective.¹⁴ I won’t go into detail about this effect because examining if the Fox and Tversky decision problem satisfies the conditions necessary to implement the recognition heuristic is the only thing I am concerned with. Only after determining if the recognition heuristic can be applied to the Fox and Tversky decision problem will such discussion become relevant.

V. Conditions For Using the Recognition Heuristic

In the Goldstein and Gigerenzer decision problem, the recognition heuristic seems to work rather well. However, it is important to recognize a pair of underlying conditions that must be met in order to implement this particular heuristic. First, the heuristic can only be applied in binary decision situations and second, only in cases where one of the two objects is not recognized. Of course, in the decision problem posed by Goldstein and Gigerenzer, these conditions were satisfied because there was a binary decision between two cities and (in the case of the German

students who were able to use the heuristic) one was recognized and the other was not.

The second condition, which I will call the “correlation condition,” is that the recognition heuristic is “domain-specific in that it only works in environments where recognition is correlated with the criterion.”¹⁵ However, this correlation is not always readily apparent; since the criterion is inaccessible to the agent making the decision, there must be a mediator that exists in the “known environment” in order to correlate the unknown criterion (in this case, population) with the recognized object (San Diego). The mediator establishes this correlation by “having the dual property of reflecting (but not revealing [directly]) the criterion and also being accessible to the senses.”¹⁶ In the case presented here, the mediator is the newspaper. There are three variables that describe mediator’s relationship between the criterion and the agent faced with the decision problem (this relationship is drawn out in fig. 1 above). These variables are the surrogate correlation, the ecological correlation, and the recognition validity. The surrogate correlation is between the mediator (which is a surrogate for the inaccessible criterion) and the recognized object. In this case, the correlation is between recognizing San Diego and the number of times the newspaper mentions it. The ecological correlation describes the relationship between the mediator and

¹⁴ *Ibid.*, 45

¹⁵ *Ibid.*, p.41.

¹⁶ *Ibid.*

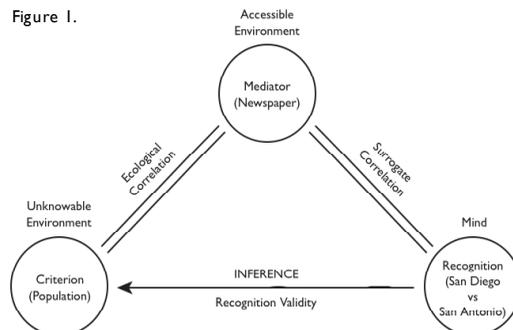
the criterion. Here, the population of San Diego is correlated with the number of times it appears in the newspaper without revealing the actual population. Lastly, the recognition validity is the proportion of “correct” answers given by use of the recognition heuristic.¹⁷ Even though the particular decision problem presented here fits the conditions for applying the recognition heuristic, this correlation condition is more complex than the first condition (of binary choice and one recognized object), and so great care must be taken in applying this heuristic to Fox and Tversky’s decision problem.

VI. Justification For Applying the Heuristic

Before looking at how this heuristic might be used in the decision problem of Fox and Tversky, it might be appropriate to consider why this particular heuristic seems appropriate to their problem. It is important to provide some justification so that my analysis amounts to more than the random application of a heuristic to an unrelated decision problem. The justification for this attempt is found in the underlying conditions involved with the Fox and Tversky decision problem. Their problem presents an agent with the choice between two objects, Bag A and Bag B, with one choice being clear due to the knowledge of probabilities and the other choice being vague as a result of unknown probabilities. Thus, the Fox and Tversky

decision problem seems close to the same types of decision problems that are compatible with the recognition heuristic,

Figure 1.



viz. a binary choice between one recognized and one unrecognized object. Without any further analysis, one might think that the recognition heuristic is applicable to the Fox and Tversky decision problem, but Goldstein and Gigerenzer state that the recognition heuristic is not a general-purpose strategy, for the necessary correlations do not hold in all domains. However, given the similarity in the decision problems described above, I think this exercise is relevant in attempting to apply this heuristic to the Fox and Tversky decision problem in order to determine how useful or adaptable such heuristics are to similar decision situations. Only by attempting experiments such as these will we be able to know the scope of the domains of relevance for these heuristics. This analysis is especially relevant because if a heuristic as simple as the recognition heuristic is unable to be applied to simple decision situations like the one presented

¹⁷ *Ibid.*, 42

by Fox and Tversky, (that is, if the domain of relevance is rather narrow) then the practical logistics of the self-proclaimed “fast and frugal” heuristics in the “adaptive toolbox” of bounded rationality will need to be questioned. Before we turn to the heuristics of bounded rationality to solve all of our problems, experiments such as the one attempted here are necessary to determine the limitations of such an approach.

VII. Application Analysis

With this background, it is now appropriate to see if the recognition heuristic can be applied to the decision problem presented by Fox and Tversky. The motivation here is to see if the recognition heuristic can be applied to describe the same decision pattern in order to account for ambiguity aversion in a theory of rational decision making that doesn't run into such problems as preference reversal. This analysis will be attempted by determining if the Fox and Tversky decision problem fits both necessary conditions for implementing the recognition heuristic and then discussing the consequences of the results. Recall the choice problem presented by Fox and Tversky: agents are to decide how much they are willing to pay for a ticket to gamble on drawing a red or a black chip from two different bags. I will assume that the “willingness to pay” as reflected in the price is indicative of which bag the agent thinks is more “likely” to produce a winning

result. Thus, a higher price given to one bag over the other reflects the agent's belief that that bag is more likely to win for them than the other. I will use this idea of the “likelihood of winning” as our criterion in the heuristic model. Hence, the decision problem becomes: which of the two bags (Bag A or Bag B) do you think is more likely to produce a winning bet?

In the Fox and Tversky experiment, one group of people chose in a comparative situation, while the other two were in non-comparative situations. Since the recognition heuristic can only be applied in binary choice decisions, it is not applicable to the latter two groups that decided in isolation. I will only be able to look at the decision problem in the comparative group, for that's the only relevant group for this analysis (since I'm not trying to account for ambiguity aversion). I am looking to see if the recognition heuristic will instruct people to choose the clear bet over the vague bet given the same decision problem.

Now, the agent is deciding between two options, “Bag A” where the bet is “clear” because the probability (of drawing a particular color chip) is known to be $\frac{1}{2}$ and “Bag B” where the bet is “vague” as a result of unknown probability. The first condition of applying the recognition heuristic is that the decision problem must be binary and that one of the two objects must be recognized while the other is not. Clearly, the agent is deciding between two objects, “Bag A” and “Bag B”, so the binary

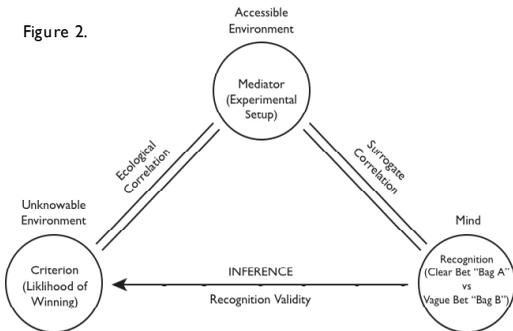
requirement holds. For the second part of this condition, we will assume that the agent recognizes “Bag A” as the clear bet by virtue of the known probability, while the agent doesn’t recognize “Bag B” as being a clear bet because the probabilities are unknown. Under this assumption, the first condition of applying the heuristic met. Satisfying the correlation condition relies on finding a suitable mediator to correlate the bags with the criterion (the likelihood of winning). In this case, the only possible mediator is the experimental setup because it completely dominates the agent’s “known environment” relative to this decision problem. The last step in determining if this decision problem satisfies the conditions for applying the recognition heuristic is if the mediator can establish surrogate and ecological correlations with the recognized object and the criterion (this decision problem is outlined in fig. 2 above). In Goldstein and Gigerenzer’s example, they used the frequency of newspaper articles mentioning the recognized object and the frequency of the object appearing in the newspaper to establish these two correlations. The problem here is slightly different. Since there aren’t multiple occurrences of “Bag A, clear bet” and “Bag B, vague bet”, we can’t use frequency to generate these correlations. However, let’s assume that this correlation is triggered by the respective probabilities of each Bag, just as the respective frequencies trigger the correlation in the newspaper case. That is,

the probabilities presented by the mediator are what triggers the recognition of Bag A being the “clear bet” and Bag B being the “vague bet”.

Given these assumptions, is the correlation condition of applying the recognition heuristic satisfied in this case? The surrogate correlation between the mediator and the object relates the recognition of the object to the agent. In the first case, the frequency of newspaper articles about the city was correlated to the number of people recognizing the city. Here, the probability corresponding to the bag is correlated with the number of people recognizing the bag as a “clear bet”. Hence, $p(\text{Bag A}) = \frac{1}{2}$, making it the “clear bet”, and $p(\text{Bag B}) = ?$ making it not recognized as the “clear bet”. Though somewhat shaky, this surrogate correlation establishes recognition of the “clear bet” and is provided by the mediator. What about the ecological correlation? This is where the application runs into problems. The mediator needs to be able to establish an ecological correlation by indirectly providing the agent with a correlation between the criterion and the recognized object. In the Goldstein and Gigerenzer example, the newspaper indirectly established the correlation between population (the criterion) and the recognized object (San Diego) by the frequency that San Diego appeared in the newspaper independent of its criterion. In our example, the experimental setup would have to correlate the likelihood of

winning with the “clear bet” in such a way that the agent thinks that $\text{Likelihood}(\text{clear bet Bag}) > \text{Likelihood}(\text{vague bet Bag})$. However, there is nothing in the mediator that indirectly (or directly for that matter) implies this because the criterion of Bag A and Bag B are dependant on their respective probabilities and $p(\text{Bag B})$ is not accounted for in the experimental setup.

Figure 2.



The failure to establish this ecological correlation makes application of the recognition heuristic impossible in this decision problem. Thus, even though the recognition heuristic seemed to be applicable and satisfied the first condition, it is unable to satisfy both necessary conditions. We must conclude that the recognition heuristic is not applicable to the Fox and Tversky decision problem.

VIII. What Went Wrong?

As stated above, this heuristic only works in a specific domain of relevance. As seen here, Fox and Tversky’s decision problem is not in this domain, though the decision problem seemed similar enough at the outset to attempt such an application.

One problem in this attempt was that the mediator was unable to establish the ecological correlation, which means that the criterion was not correlated with recognition. This ecological correlation is impossible for two reasons. First, the criterion of the “clear bet” (likelihood of winning) is not independent of the unrecognized object’s criterion since the relative probabilities of each bag can determine the likelihood of winning a bet on either Bag A or Bag B. In the Goldstein and Gigerenzer example, it would be like claiming the population of San Diego depends on the population of San Antonio and vice versa. Second, even if we disregard this dependence, the mediator failed to establish an ecological correlation because it didn’t provide the agent with the $p(\text{Bag B})$ in order to correlate likelihood of winning with the recognized bag. Of course, if this were the case, then the condition of recognition is also moot and the problem doesn’t satisfy the first criterion either. Since the Fox and Tversky choice problem’s “knowable environment” was confined to the experimental setup, there is nothing else that can act as mediator here. I think it’s important to realize this dependant relationship between the correlation and the mediator because it significantly restricts the domain of relevance for the heuristic.

A more fundamental problem here has to do with limitations on how robust this heuristic is. There are two aspects that limit the domain of relevancy; the definition

of “recognition” and the simplicity of correlation. It appears that the definition of “recognition” is more restrictive than the authors let on in distinguishing between the novel and previously experienced. Recognition for the heuristic is merely recognizing an object A based on prior encounters with A. In our example, the recognition assumption was applied to a specific property of the bags viz. being a “clear bet” or a “vague bet,” rather than the bags themselves. This property of the bag must be information about the object, whereas recognition only applies to the object itself. “Ignorance” in the notion of “ignorance making us smart” is restricted to instances of recognition. However, ignorance must be more than just not recognizing something, as illustrated in this experiment; it also includes situations where information is lacking. It’s apparent that we cannot equate knowledge, or lack of information, with the restricted notion of causal recognition and non-recognition. This concept of recognition used in the heuristic seems too limiting even when attempting to apply it in this simple case.

The second fundamental problem of robustness here is that the recognition heuristic relies on a single recognized object A and one particular attribute X of that object. Even though the Fox and Tversky decision problem was simple and binary, it involves more than the simple correlated relationship between an object A and a specific attribute X. It also involves the interaction of two related objects that

are dependant on one another, while the recognition heuristic assumes that the objects under consideration and their criterion are completely independent of each other. Thus, it appears that the domain of relevance of the recognition heuristic is more restricted than was anticipated, as illustrated by its complete failure to give an accounting for the decision problem in Fox and Tversky.

IX. Conclusion

This heuristic is one of the many tools that have been developed in the decision theory known as bounded rationality. The reason that this approach to decision making is interesting is that it doesn’t rely on the traditional method of maximizing expected utility. Therefore, it’s immune to problems such as preference reversal. However, before we view this approach as a panacea for the traditional theory, we must closely analyze these heuristics to determine how relevant they are in certain situations. Determining the relevance of every heuristic out there is beyond the scope of this paper. However, as in the case presented here, we might have good reason to attempt just such a thing. Although, if simple heuristics such as the recognition heuristic fail to apply easily to other simple situations, then their domains are rather restricted. Though the authors admit this, it brings into doubt the general usefulness of this particular heuristic in the so-called “adaptive toolbox”.

At this point the question then becomes, 'is the adaptive toolbox offered as a supplement to traditional expected utility theory, or a replacement of it?' If one wants to claim the latter, then the case presented here is a problematic example, because in order give a complete theory of decision making with the "adaptive toolbox" model, there seems to be a logistical problem in implementing and choosing which heuristics have domains relevant to different choice situations. To adopt the imagery of the "adaptive toolbox", it's like needing to carry around an infinite number of screwdrivers for infinite different sized screws. Eventually the toolbox gets so big that it becomes too cumbersome for practical use. We may then need to revert back to the "hammer" of tradition expected utility theory and bang everything out, even if it may not be the ideal tool for the job. Hence, the heuristic approach needs to be supplementary to traditional expected utility theory. Hopefully different heuristics can be developed that can account for situations where the traditional approach fails. As is clearly the case here, the recognition heuristic isn't a satisfactory approach to explain decision patterns exhibiting preference reversal and ambiguity aversion in the context of Fox and Tversky's experiment; so another method must be sought after.

Works Cited

- Fox, Craig R., and Amos Tversky.
"Ambiguity Aversion and Comparative Ignorance." *The Quarterly Journal of Economics* 1220, no. 3 (Aug., 1995): 585-603.
- Gigerenzer, Gerd, and Daniel Goldstein.
"The Recognition Heuristic: How Ignorance Makes Us Smart." *Simple Heuristics That Make Us Smart*. Edited by Gerd Gigerenzer, Peter M. Todd, and The ABC Research Group. New York: Oxford University Press, 1999.
- Simon, H.A. "Rational Choice and the Structure of Environments," *Psychology Review* 63 (1956): 129-138.