

Iterated Random Selection as Intermediate Between Risk and Uncertainty

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

Jeffrey Helzner
Columbia University
jh2239@columbia.edu

Abstract

In [7] Hertwig et al. draw a distinction between *decisions from experience* and *decisions from description*. In a decision from experience an agent does not have a summary description of the possible outcomes or their likelihoods. A career choice, deciding whether to back up a computer hard drive, cross a busy street, etc., are typical examples of decisions from experience. In such decisions agents can rely only of their encounters with the corresponding prospects. By contrast, an agent furnished with information sources such as drug-package inserts or mutual-fund brochures—all of which describe risky prospects—will often make decisions from description.

In [7] it is shown (empirically) that decisions from experience and decisions from description can lead to dramatically different choice behavior. Most of these results (summarized and analyzed in [6]) are concerned with the role of risk in decision making. This article presents some preliminary results concerning the role of *uncertainty* in decision making. We focus on Ellsberg's two-color problem and consider a chance setup based on double sampling. We report empirical results which indicate that decisions from description where subjects select between a clear urn, the chance setup based on *double sampling* and Ellsberg's *vague* urn, are such that subjects perceive the chance setup at least as an intermediate option between clear and vague choices (and there is evidence indicating that the double sampling chance setup is seen as operationally indistinguishable from the vague urn). We then suggest how the iterated chance setup can be used in order to study decisions from experience in the case of uncertainty.

Keywords. decisions from description, decisions from experience, random selection, uncertainty

1 Introduction

Consider a scenario in which a well-educated couple must decide whether or not their child should receive a particular vaccination. To assist in their decision making, the couple is provided with statistics concerning the frequency of serious, adverse reactions associated with the vaccination in question. Though the frequency of such adverse reactions is quite low, the couple is reluctant to have their child vaccinated. Concerned, the child's pediatrician provides reasons in favor of vaccination, noting that she herself in fact had chosen to vaccinate her own children. What might explain the difference between the judgement of the child's parents and that of the child's pediatrician? One plausible explanation that is of general, theoretical interest focuses on the way in which the relevant parties are acquainted with the chances associated with an adverse reaction. While the child's parents are provided with frequencies, and presumably the child's pediatrician is privy to this information, the pediatrician can also draw upon her own clinical experience.

Hertwig et al. propose to analyze cases of this type in terms of a distinction between *decisions from description* and *decisions from experience*. As suggested by the scenario in the previous paragraph, which is among the scenarios that provide motivation for the work reported in [7], decisions from description are made with the benefit of a description of the relevant chance mechanism, e.g., associated probabilities, while decisions from experience are informed by repeated encounters with the chance mechanism itself, i.e., sampling. As noted in [7], the vast majority of experimental work on decision making has focused on decisions from description. The following example, taken from Kahneman and Tversky's 1979 classic on prospect theory [8], illustrates the methodology typical among work in this area:

Example 1. *Which of the following do you prefer?*

Alternative 1 pays \$5000 with probability .001 and \$0 with probability .999. Alternative 2 pays \$5 with probability 1.

Is there any reason to think that this focus on decisions from description has been significant with respect to the results gathered through numerous studies which employ the sort of methodology illustrated in Example 1? In [7], Hertwig et al. answer this question in the affirmative. Specifically, they present evidence indicating that people tend to “underweight the probability of rare events” when making decisions from experience. This is in stark contrast to the well-known results of Kahneman and Tversky, based on items such as Example 1, which indicate that people tend to overweight the probability of rare events when making decisions from description. Thus, it seems that Hertwig et al. have isolated an important psychological effect. The opening scenario clearly illustrates the effect Hertwig et al. have isolated. The child’s parents are presented with frequencies which they interpret as a description of the relevant chance mechanism. As predicted by Hertwig et al., the child’s parents overweight the probability of an adverse reaction and decide not to vaccinate. By contrast, the pediatrician’s decision making is informed by her clinical experience. As predicted by Hertwig et al., the child’s physician underweights the probability of an adverse reaction and recommends vaccination.

Hertwig et al. restrict their attention to decision making under risk. In particular, the descriptions that they employ include a numerically precise probability distribution. The main purpose of this paper is to begin an investigation into the possibility of a *description-versus-experience* effect in the context of decision making under uncertainty, i.e., in contexts where the description of the relevant chance setup does not determine a numerically precise probability distribution. On the basis of experimental evidence reported in this paper we conjecture that the gap between *vague* and *clear* is less pronounced in the case of decisions from experience than in the case of decisions from description.

2 From Risk to Uncertainty

Consider the design of the first study reported by Hertwig et al. in [7]. The subjects were divided into two groups: Description and Experience. Those in Description were presented with several choice problems, each of which consisted of a pair of risky alternatives. For example, one such choice problem consisted of the alternatives $(4, .8)$ and $(3, .1)$, where (m, p) denotes the risky alternative that pays amount m with probability p and pays amount 0 with probability $1 - p$.

Let (m, p) be a risky alternative of the indicated sort. One can construct a chance setup that satisfies description (m, p) . For example, a chance setup for the alternative $(4, .8)$ could use random draws (with replacement) from an urn consisting of 80 black balls and 20 white balls. The implementation of the second group, Experience, is less familiar. Consider a decision-from-experience counterpart to the choice problem consisting of $(4, .8)$ and $(3, .1)$. One could, for example, present the subject with two buttons, say A and B , where pressing A (B) results in a trial on a particular chance setup corresponding to $(4, .8)$ ($(3, .1)$). The subject, who sees the result of each trial (e.g., in the case of A , whether the payoff would have been 4 or 0), is permitted to sample the two chance setups as many times as they wish before they are required to make a choice and play one of the two setups for real payoffs. Essentially, this is the way in which Hertwig et al. study decision making from experience.

What happens when we move from risk to uncertainty (where probabilities are imprecise)? The obvious candidate on the description side is familiar through the presentation of Ellsberg problems such as following:

Example 2 (Ellsberg’s two-color problem [4]). *Consider the following two cases:*

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.*

Consider now the following questions: How much would you be willing to pay for a ticket that pays \$55 (\$0) if the next random selection from Urn A results in black (white) ball? Repeat then the same question for Urn B.

Following the above presentation, an uncertain alternative over a pair of prizes (only one of which is nonzero) can be specified by providing the amount of the nonzero prize and the *set* of probabilities that are associated with that prize. Thus, for example, $(55, \{\frac{i}{100} \mid 0 \leq i \leq 100\})$ is the uncertain alternative in which the probability of winning \$55 is known to be in $\{\frac{i}{100} \mid 0 \leq i \leq 100\}$.

Presenting alternatives in this way has the virtue of generalizing the risky alternatives that are employed by Hertwig et al., since these risky alternatives are simply those of the form $(m, \{p\})$. However, once the probabilities are allowed to be indeterminate, it is not clear how to complete the analogy in a way that would support decision from *experience* under uncertainty. Recall the desired relationship in the case of

risk. Given a risky alternative, e.g., $(m, \{p\})$, one can construct a corresponding chance setup, i.e., one that satisfies description $(m, \{p\})$. Doing the analogous thing in the case of uncertainty would seem to require chance setups that implement indeterminate probabilities. Are there such things in any interesting sense? After all, the uncertainty described in Ellsberg-type examples is purely epistemic, e.g., the ratio between black ball and white balls in Urn B *is* determinate even though it is not known to the decision maker. On the other hand, consider the following description of a chance setup:

B^* : First, select an integer between 0 and 100 at random, and let n be the result of this selection. Second, make a random selection from an urn consisting of exactly 100 balls, where n of these balls are solid black and $100 - n$ are solid white.

As in the case of Urn B, the outcome of a trial on chance setup B^* depends on a random selection from an urn such that the ratio of black balls to white balls is not known to the subject. However, unlike the case of Urn B, the subject knows that the urn that is sampled in the second stage of B^* is determined by a random selection in the first stage of B^* . According to at least one familiar line of reasoning, this second consideration suggests that a play on B^* is equivalent to a play on Urn A, rather than Urn B. The indicated line of reasoning is roughly as follows: The random selection in the first stage entails that, for each integer i , where $0 \leq i \leq 100$, there is a probability of $\frac{1}{101}$ that the urn sampled in the second stage consist of i black balls and $100 - i$ white balls. Moreover, according to this line of reasoning the random selection in the second stage entails that if i is selected in the first stage, then the probability of selecting a black ball in the second stage is $\frac{i}{100}$. This line of reasoning then continues by combining the first and second stage probabilities to conclude that the probability of getting a black ball on a trial of B^* is $\frac{1}{101}(\sum_{i=0}^{100} \frac{i}{100}) = \frac{1}{2}$, as in the case of Urn A. There are, of course, well-known responses to this line, the most obvious being one that questions the relevance of a chance setup's long-run behavior when it comes to assigning probabilities for a single trial of the setup; here we are assuming that the relevant probabilities are based on frequencies rather than something like propensities.

A less familiar response maintains that one has complete uncertainty over the collection of chance setups that satisfy description B^* and that, since some of the setups will select a black ball on their next trial while others will select a white ball, one has complete uncertainty with respect to the outcome of the next trial. For example, even if one maintains that there

are truly nondeterministic chance setups, deterministic chance setups are common and are of the sort that Hertwig et al. employ in [7]. If random selection is understood to mean selection by a mechanism such that (1) future behavior of the mechanism cannot be predicted from a mere knowledge of its past behavior and (2) the various possible outcomes are distributed evenly in the long run – and these are important matters that will be considered in the sections that follow – then the description of B^* is compatible with the use of such deterministic mechanisms.

For all the subject knows, the first stage selection in B^* can be made according to a deterministic process that will select 33 on its next run, while the second stage mechanism will be made according to a deterministic mechanism that will select a black ball on its next draw from the 33:77 urn. Similarly, it is compatible with the information that is presented to the subject that the first stage selection in B^* will be made according to a deterministic process that will select 61 on its next run, while the second stage mechanism will be made according to a deterministic mechanism that will select a white ball on its next draw from the 61:39 urn. The subject has complete uncertainty over these selection mechanisms and, more generally, over the collection of all chance setups that might be used to carry out the selections in B^* . Since the final outcome, i.e., the selection of a black or white ball, is a function of the chance setups which are employed, at least where deterministic mechanisms are used, complete uncertainty over the collection of these mechanisms suggests complete uncertainty with respect to the final outcome. Note that this line is rather extreme, since it suggests complete uncertainty even in situations where the second stage urn is fixed, e.g., as in a chance setup that makes selections with replacement from Urn A. Rather than trying to achieve consensus with respect to such *a priori* considerations, we now turn our attention to psychological matters.

3 Study

What is the psychological relationship between the description of Urn B and B^* ? To investigate this question we asked subjects to state their maximum buying prices with respect to hypothetical situations involving the descriptions at issue. Our study included 89 undergraduates from Carnegie Mellon. At the time of the study the participating students were enrolled in 80-100, an introductory philosophy course at Carnegie Mellon. Each subject was presented with a questionnaire, the contents of which will now be described.

After instructing the subjects that they would be presented with questions involving hypothetical scenar-

ios, the questionnaire continued with the following tutorial on chance setups:

Think of a roulette wheel of the sort that one would find in any American casino. The casino employee spins the wheel in one direction and then sends a ball in the other direction along a track that goes around the circumference of the wheel. Eventually the ball comes to rest in one of the wheel's 38 pockets. Players expect that this setup is fair in the sense that the following conditions are satisfied: (1) In the long run the number of times that the ball lands in a particular pocket is equal to the number of times that the ball lands in any other particular pocket and (2) One cannot predict where the ball will land on the next spin simply by knowing where the ball landed on previous spins.

Roulette wheels are a special case of a more general class of systems. More generally, a *chance setup* is a system that includes a finite set $\{o_1, \dots, o_n\}$ of possible outcomes and that outputs one of these outcomes each time that it is run. Such a chance setup is *fair* just in case the following conditions are satisfied: (1) If the system were run repeatedly, then in the long run the number of trials that would result in outcome o_i would be equal to the number of trials that would result in outcome o_j and (2) One cannot predict which outcome will result from the next run of the system simply by knowing the outcome of each of the previous runs of the system.

The questionnaire then continued by instructing the subjects that “random selection” in the context of the questionnaire is to be understood as selection via a fair chance setup. This instruction was followed by three questions:

1. An urn has been filled with exactly 100 balls. 50 of the balls are black and the remaining 50 are white. A random selection from the urn will be made. **What is the most that you would be willing to spend on a ticket that pays \$55 if the random selection results in a black ball and pays \$0 if the random selection results in a white ball?**
2. Consider the following two-stage process: (1) A random selection is made from the collection $\{0, 1, \dots, 100\}$ and (2) A second random selection is made from an urn that contains exactly n black balls and $100 - n$ white balls, where n is

the result of the random selection from the first stage. Thus, for example, if 23 was the result of the random selection in the first stage, then the random selection in the second stage would be from an urn containing exactly 23 black balls and 77 white balls. **What is the most that you would be willing to spend on a ticket that pays \$55 if the random selection in the second stage results in a black ball and pays \$0 if the random selection in the second stage results in a white ball?**

3. An urn has been filled with exactly 100 balls. Each ball in the urn is either black or white. However, the ratio of black balls to white balls in the urn is unknown. A random selection from the urn will be made. **What is the most that you would be willing to spend on a ticket that pays \$55 if the random selection results in a black ball and pays \$0 if the random selection results in a white ball?**

Note that the ticket is played against the “clear” chance setup in the first question, against the “double sampling” chance setup in the second question, and against the “vague” chance setup in the third question. Thus, referring back to Example 2, the first of these questions asks the subject to price the ticket on Urn A, while the second asks the subject to price the ticket with respect to B^* , and the third asks the subject to price the ticket on Urn B.

Recognizing that the order in which the questions appeared might affect the responses, we created three different versions of the questionnaire: CDV, VDC, and DVC. Version CDV was administered to 41 subjects and presented the questions in the order given above, i.e., the question concerning the clear setup followed by the question concerning the double-sampling setup followed by the question concerning the vague setup. Version VDC, which was administered to 32 subjects, reversed this order. Version DVC was given to 16 subjects and had the question about double-sampling occurring first, the question about the vague setup occurring second and the question about the clear chance setup occurring third. In each version, subjects were instructed to answer the questions in the order that they were presented.

For each of the three groups, Table 1 shows the mean maximum buying price for the three questions. Thus, for example, the first row of the second column indicates that, in the case of the ticket on the clear chance setup, \$22.68 was the mean maximum buying price for the group of subjects that received the VDC version of the questionnaire.

Question	CDV mean	VDC mean	DVC mean
Clear	22.89	22.68	19.02
Double	14.68	9.77	7.70
Vague	5.82	7.10	3.25

Table 1

While the order of the questions appears to have had some bearing, the basic pattern of Vague < Double < Clear for the mean maximum buying prices seems robust across the three versions of the questionnaire. The mean maximum buying prices over all subjects are shown in Table 2.

Question	Mean
Clear	22.12
Double	11.66
Vague	5.82

Table 2

If we turn our attention to the level of individual subjects, then the above pattern of strict inequalities is less pronounced since approximately $\frac{1}{3}$ of the subjects gave the same maximum buying price for Double and Vague; moreover, these subjects were not distributed evenly across the three groups. However, as shown in Table 3, a clear pattern emerges if we weaken the first inequality.

Group	# $V \leq D < C$	% $V \leq D < C$
CDV	29	71%
VDC	24	75%
DVC	12	75%
All	65	73%

Table 3

The first column of Table 3 shows the number of subjects in each group (and overall) that satisfied the Vague \leq Double < Clear pattern, while the second column shows the associate percentages. As Table 3 shows, these percentages are quite stable across the three groups individually and their union. It is also worth noting that, as shown in Table 4, relatively few of the subjects gave the same maximum buying price for Clear and Double.

Group	# $C = D$	% $C = D$
CDV	12	29%
VDC	6	19%
DVC	4	25%
All	22	25%

Table 4

The first column of Table 4 shows the number of subjects in each group (and overall) that gave the same maximum buying price for Clear and Double, while the second column shows the associate percentages.

The data also reveal as well an interesting result if we consider the values for Vague, Double, and Clear given by the mean maximum buying prices for VDC, DVC and CDV, respectively. These values are important because these are the results for the three cases without considering comparisons – subjects were instructed to answer the questions sequentially and not to return to previous questions. For example, since Vague occurs first in VDC, it seems reasonable to assume that it is evaluated in a non-comparative context. Table 5 shows the mean maximum buying prices for these three non-comparative cases:

Question	Mean
Clear	22.89
Double	7.70
Vague	7.10

Table 5

These results suggest an operational identification of Vague and Double, which have almost identical values; this is of interest to us since Double, unlike Vague, seems to be implementable in a way that could support decisions from experience, and this is something that we will revisit in the following section. Furthermore, these values for Vague and Double are clearly separated from the mean maximum buying price for Clear as reported in Table 5. Perhaps the only concern here is that the number of subjects who received VDC is low in comparison to the number of subjects in the other two groups. We expect to run a larger experiment including decisions from experience. In this case it would be interesting to see whether the pattern verified in the pilot is robustly maintained.

Another issue that we intend to address in future experiments concerns the worry, expressed by one anonymous referee, that the complexity of B^* rather than its status with respect to uncertainty is responsible for its lower price. This raises two interesting issues. First, how might we control for this in future experiments? Perhaps one way to do this would be to have the subjects explain the reasoning behind their responses. We can certainly ask the subjects why they priced one alternative lower than another. The second issue concerns the relationship between complexity and uncertainty. In the present context this seems to line up with the familiar distinction between imprecision and indeterminacy. One might claim that the credal probabilities in the case of B^* are imprecise

but not indeterminate, e.g. that the rational agent is committed to a particular credal probability but is unable to identify that particular distribution. This sort of situation is not unlike the imprecision that arises in connection with the measurement of physical concepts, e.g. length or weight. In contrast, one might claim that in the case of Urn B the rational agent's credal probability itself – rather than just its estimation of that credal probability – ought to be indeterminate. Surely the distinction is not mere stipulation, but what is wrong with maintaining that the rational agent's credal probabilities with respect to B^* should also be indeterminate (i.e. that the rational agent is not committed to a determinate credal probability in such a case)? It seems that considerations of this sort lead in the direction of bounded rationality, in particular the tenability of capacity independent notions of rationality.

4 Discussion

The results of our study suggest that double sampling is perceived as something in between risk and uncertainty when comparative contexts are allowed, but is there a way to make this intermediate position more transparent? Perhaps one way to do this would be to describe a mixed chance setup in which the subject is told that the selection will be made from Urn A with probability p and from Urn B with probability $1 - p$. One could then attempt to identify the value of p at which the subject's maximum buying price is equal to the maximum buying price that the subject stated with respect to B^* . The value of p so identified could be taken as a representation of the intermediate position that B^* occupies between risk (Urn A) and uncertainty (Urn B).

One drawback to using descriptions of mixed chance setups in the manner suggested above is that such descriptions do not appear to fit, at least psychologically, into the sets-of-probabilities approach to representing credal states. While the set of all distributions p such that $p(\text{Black}) = \lambda p_1(\text{Black}) + (1 - \lambda)p_2(\text{Black})$, where $p_1 \in X$ and $p_2 \in Y$ might seem like a natural representation of a mixture with probability λ on X and probability $1 - \lambda$ on Y , some preliminary results reported in [2] suggest that this is not the case.

If descriptions of double sampling are perceived as something distinct from risk, what might an implementation of such a description look like in a study of decision making from experience? Recall the study by Hertwig et al. in [7]. They implemented the description of a risky alternative, such as (m, p) , as an appropriate chance setup. A trial on this chance setup is activated by a button on a computer screen. After

pushing this button, the subject sees the outcome of the trial on the computer screen. If we attempt to implement B^* in such a way that the subject sees only the final result of the two-stage process after pushing the appropriate button, then we run into problems. The issue is that there is nothing to guarantee that such an implementation of B^* could not just as well serve as an implementation of Urn A. By assumption, a chance setup that satisfies B^* will, in the long run, draw j in the first stage approximately $\frac{1}{101}$ th of the time. Moreover, in the long run, trials in which j is drawn in the first stage result in the selection of a black ball approximately $\frac{j}{100}$ th of the time. Hence, in the long run, $(\frac{1}{101})(\frac{j}{100})$ th of the trials result in a black ball drawn from the urn having j black balls and $100 - j$ white balls. Thus, in the long run, approximately

$$\left(\frac{1}{101}\right)\left(\frac{1}{100}\right)\sum_{j=0}^{100} j = \frac{5050}{10100} = \frac{1}{2}$$

of the trials result in the draw of a black ball, which agrees with the limiting frequencies for an implementation of Urn A. If we assume that such an implementation of B^* would also satisfy the condition that future behavior cannot be predicted solely from a knowledge of past behavior, and this seems to be a psychological matter, then it appears that there is nothing to prevent such an implementation of B^* from serving as an implementation of Urn A. Clearly the implementations of B^* and Urn A must be distinct in a meaningful way if one is to conduct the desired study of decision making from experience. We now consider other proposals for implementing B^* .

One way to avoid the sort of problematic collapse discussed at the end of the previous paragraph would be to make both stages of the double sampling visible to the subject. Thus, for example, pressing the appropriate button on a computer screen for the first time would run the first-stage selection, and the result of that selection would be shown to the subject, e.g., that the urn with 35 black balls and 65 white balls had been selected. Pressing the button for the second time would initiate a draw from the urn that had been selected, and the result of that second-stage selection would then be shown to the subject, e.g., that a white ball had been drawn. Making both stages of double sampling visible to the subject avoids the problematic collapse since such an implementation of B^* no longer qualifies as an implementation of Urn A.

However, there is a possible objection to this design, based on the fact that it implements a sort of hybrid experimental condition that does not correspond purely to decisions from experience or decisions from description. Typically in decisions from experience

the subjects do not have access to the probabilities of the option considered. In the previous design one makes at least intermediate (i.e., first stage) probabilities explicit by revealing the composition of the selected urn.

There is a remedy to the previous objection via the implementation of the following experimental design. This design assumes that the following four buttons are available to the subject: PLAY, SELECT GAME, V, and C. The subject's initial choice concerns V and C. Button C implements the "clear" scenario that was presented in the questionnaire. If C is selected, then the agent can press PLAY repeatedly. Pressing PLAY samples from the implementation of the clear urn. While the urn structure is hidden from the subject, the subject sees the associated payoffs, \$55 if black and \$0 if white, after each pressing of PLAY. If V is selected, then the subject is instructed to press SELECT GAME. Unbeknownst to the subject, pressing SELECT GAME selects an implementation based on one of the 101 possible urns considered above (i.e., an urn consisting of n black balls and $100 - n$ white balls for some $n \leq 100$).

The following algorithm is used to select a game: consider the space of all possible ordered sequences of 101 urns. Then a sequence in this space is selected at random and fixed. When the game starts and the agent presses SELECT GAME for the first time the first urn in the sequence is selected. Say that the subject has pressed SELECT GAME n times. Then when he presses SELECT GAME once more the selection mechanism picks the urn in the $n + 1$ position in the sequence and samples it every time that PLAY is selected. Now at each point the probability of white or black will depend on the previous actions of the subject playing the game. Since probabilities should not be attributed to acts these probabilities remain indeterminate. When the sequence terminates the algorithm starts again at the initial point of the selected sequence.

It is important to remark that what *does not* have a determinate probability is the color of the first ball prior to selecting or not a game (i.e. prior to choosing to play). Likewise for the probability of the n th ball prior at the moment of the choice whether or not to select a game for the n th time. After the agent selects a game (i.e. after the agent decides to play the game) we have a uniform precise probability over the 101 configurations of the urn.

In addition after selecting games a few times and sampling them the calculation of the probabilities of the color of the ball in the next trial grows ever more complicated after conditioning on what has been done and

on what has been seen from past plays of the game. Even for an ideal agent these probabilities will be imprecise. The agent will have bounds of the values of the probabilities or he can form qualitative judgments comparing probabilities but the agent will not have precise probabilities at his disposal. So, under a normative point of view there is indeterminacy at the moment contemporary to the selection of the game, and after a few trials there will be imprecision in the corresponding probabilities.

When we consider bounded agents the situation is even worse. The agent might not remember well what he did in the past and what he saw in the past and we can have recency effects as well. So, in the real situation we have to deal with imprecise probabilities and choices under uncertainty.

After SELECT GAME is pressed the subject has a choice: explore the game that was selected by pressing PLAY or select a (possibly) new game by pressing SELECT GAME. Pressing PLAY samples from the current game. While the urn structure of the game is hidden from the subject, the subject sees the associated payoffs, \$55 if black and \$0 if white, after each pressing of PLAY. The agent can interact with the two buttons as long as he wants. Notice that it is perfectly possible that an agent selects a game and then presses PLAY repeatedly without selecting any other game. Notice that in this case the argument in terms of frequencies fails (the agent does not see data from all urns, he just considers a single (or a few) urns). So, the shift to the design where the two stages are visible is essential in order to avoid the argument for the collapse into urn A.

Example 3. A possible session involving V:

The subject first presses V.

On screen: *Select a game by pressing SELECT GAME. OR EXIT*

The subject presses SELECT GAME.

On screen: *You have been awarded a game. You can play this game by pressing PLAY.*

The subject presses PLAY and a payoff appears, for example:

On screen: *You won \$55.*

The subject then agent faces the choice of pressing PLAY again or pressing SELECT GAME or EXIT (in which case he faces the election of V and C again).

If the agent chooses instead button C, he will have the option of pressing PLAY as many times as he wants. A payoff will appear each time that PLAY is pressed. At any time he can STOP and choose between V and

C.

Example 4. A possible session involving C:

The subject first presses C.

On screen: *Press PLAY.*

The subject presses PLAY and a payoff appears, for example:

On screen: *You won \$0.*

The subject will can press PLAY as many times as he wants. A payoff will appear after each time that PLAY is pressed.

After receiving feedback from these two buttons the agent has to select V or C and in this case he will play for real money. Of course, if he selects V a new game will be selected by pressing SELECT GAME and he will receive the payoff determined by the next pressing of PLAY, i.e., by sampling the urn corresponding to the current game.

This design makes visible the two-stage nature of the V button but, as in the case of decisions from experience for risk, the agent does not receive any information about intermediate probabilities. Notice that the algorithm used in the proposed implementation of decisions from experience is a particular instance of a selection via a fair chance set up (as described to subjects in the tutorial). Since it is clear that in this case there is no collapse of the implemented mechanism with the clear urn, it follows that in general there is no reason to expect a collapse of B^* and C. This shows that the operational identification of the V and B^* conditions might not just be attributable to a statistical error on the part of the subjects.

We conjecture that this design of decisions from experience will also avoid an identification of the V and C conditions. We also conjecture nevertheless that the gap between the V and C conditions (buttons in decisions from experience) will not be as severe between the gap between the corresponding “vague” condition and the “clear” conditions in the case of decisions from description. This is because it is unlikely that the subject encounters rare events while obtaining feedback by interacting with button V, and this suggests that the subject will remain ignorant of their existence (examples of extreme values or rare events will be the case where either Black or White are zero or very low in the sampled urns).

We conjecture therefore that this proposal will show a significant difference between decisions from description and decisions from experience, demonstrating that the distinction between these two types of decisions is robust and applicable not only to risk but also to the case of uncertainty.

5 Further Considerations

We conclude by mentioning another issue that is raised by our study, an issue that seems to have general significance for experimental work on decision making. An unusual aspect of the questionnaire that we used in the study that is reported in Section 3 is the fact that it is explicit about what is meant by a random selection. While references to random selection are common in experimental work on decision making, these references are seldom accompanied by something like the tutorial on fair chance setups that was part of our study. It is natural to wonder if this makes a difference. While we have yet to conduct a study of this particular question, we do have data from an earlier study that seems to suggest that it does make a difference if one is explicit about what is meant by a random selection.

As part of the study that we reported in [1], we used a questionnaire that asked subjects to state their maximum buying price for what were essentially questions Clear and Vague as presented in Section 3. It is important to note that the questionnaire that was used in [1] did not include any tutorial on fair chance setups, nor for that matter did it include any elaboration regarding the nature of random selection. Finally, it should be noted that the subjects in this earlier study were, like those of the study reported in Section 3, undergraduates at Carnegie Mellon who, at the time of the study, were enrolled in 80-100, which as noted in Section 3 is an introductory philosophy course. Table 6 shows the mean maximum buying prices for the two groups in the earlier study that received Clear as the first question on their questionnaire.¹

Group	Mean for Clear (2005)
I	15.33
II	13.65

Table 6

These values seem significantly less than the mean maximum buying prices for Clear that are reported in Section 3. These differences seem striking when one considers the mean maximum buying prices that were obtained for Vague in the earlier study. Table 7 shows the mean maximum buying prices for the two groups in the earlier study that received Vague as the

¹The two groups were distinguished by the fact that they were given slightly different questionnaires. Both groups received a questionnaire that had Clear as the first question, but there were some differences between the two questionnaires in their later sections. We do not think that these differences are significant in the present context, but the interested reader can consult [1] for a detailed description of the questionnaires that were involved.

first question on their questionnaire. ²

Group	Mean for Vague (2005)
I	5.42
II	6.4

Table 7

These values seem more in line with the mean maximum buying prices for Vague that are reported in Section 3. As a measure of this effect, Table 8 shows the ratio of the mean maximum buying price for Vague to that of Clear.

Group	Vague/Clear
2005	.41
2008	.31

Table 8

Column 2 of the first row in Table 8 shows the average of the two values reported in Table 7 divided by the average of the two values reported in Table 6. The second row of Table 8 shows the mean maximum buying price for Vague as reported in Table 5 divided by the mean maximum buying price for Clear as reported in Table 5. Taken together, these further considerations would seem to raise an important question as to how subjects are interpreting references to random selection in those studies that do not elaborate on what is meant by such a thing.

ACKNOWLEDGMENTS:

A first version of this paper was read at the Seminar of Games and Decisions at CMU. We received numerous comments then. We want to thank in particular Teddy Seidenfeld who proposed a variant of the design of decisions from experience that we are adopting here. Isaac Levi and Paul Pedersen also provided useful comments and criticism. Finally we want to thank specially Ralph Hertwig who provided invaluable feedback and criticism. Various insightful conversations with Ralph during his recent visit to Columbia University motivated us to write this paper.

References

[1] Arló-Costa, H. & Helzner, J. (2005). Comparative Ignorance and the Ellsberg Phenomenon. *Proceedings of the Fourth International Symposium on Imprecise Probabilities and Their Applications*, eds. Cozman, F, Nau, R. & Seidenfeld, T.

[2] Arló-Costa, H. & Helzner, J. (2007). On the Explanatory Value of Indeterminate Probabilities. *Proceedings of the Fifth International Symposium on Imprecise Probabilities and Their Applications*, eds. De Cooman, G, Vejnarova, J. & Zaffalon, M.

[3] Chow, C.C. & Sarin, R.K. (2001). Comparative Ignorance and the Ellsberg paradox. *Journal of Risk and Uncertainty*, Vol. 22, No. 2, 129-139.

[4] Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, Vol. 75, No. 4, 643-669.

[5] Fox, C.R., & Tversky, A. (1991). Ambiguity Aversion and Comparative Ignorance. *The Quarterly Journal of Economics*, Vol. 110, No. 3, 585-603.

[6] Hertwig, R. (2009) The Psychology and Rationality of Decisions from Experience, *Synthese*, forthcoming.

[7] Hertwig, R., Barron, G, Weber, E.U., & Erev, I. (2003). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, Vol. 15, No. 8, 534-539.

[8] Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, Vol. 47, 263-291.

²The distinction between the two groups in this case is completely analogous to the distinction between the two groups in Table 6.