Games and Economic Behavior 79 (2013) 44–55

Contents lists available at SciVerse ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Shunning uncertainty: The neglect of learning opportunities ☆

Stefan T. Trautmann[a], Richard J. Zeckhauser[b,*]

[a] Tilburg University, PO Box 90153, 5000 LE Tilburg, The Netherlands
[b] Harvard University, Kennedy School of Government, 79 JFK St., Cambridge, MA 02138, USA

ARTICLE INFO

Article history:
Received 29 November 2010
Available online 16 January 2013

JEL classification:
D81
D83
G11
C91

Keywords:
Risk
Uncertainty
Learning
Information
Gambler's fallacy

ABSTRACT

Financial, managerial, and medical decisions often involve alternatives whose possible outcomes have uncertain probabilities. In contrast to alternatives whose probabilities are known, these uncertain alternatives offer the benefits of learning. In repeat-choice situations, such learning brings value. If probabilities appear favorable (unfavorable), a choice can be repeated (avoided). In a series of experiments involving bets on the colors of poker chips drawn from bags, decision makers often prove to be blind to the learning opportunities offered by uncertain probabilities. They forgo significant expected payoffs when they shun uncertain alternatives in favor of known ones. Worse, when information is revealed, many make choices contrary to learning. Priming with optimal strategies offers little improvement. Such decision makers violate identified requirements for making rational decisions.

Financial, managerial, and medical decisions often involve choice alternatives whose outcomes have uncertain probabilities. In contrast to alternatives with known probabilities, these uncertain alternatives offer the opportunity for learning in repeat choices. If the probability of a desirable outcome turns out low, the alternative can be avoided in future choices. If the probability turns out high, the alternative can be chosen again on future occasions. Thus if repetition is possible, the opportunity for learning brings value, making uncertain alternatives more valuable than known-risk alternatives offering identical expected value on a single trial. To gather information regarding uncertain probabilities, people typically need to actively engage uncertain alternatives, something they are reluctant to do.

A range of theoretical analyses show the benefits of learning opportunities offered by uncertain alternatives. For instance, Grossman et al. (1977) show that, if learning is possible, consumers may buy a drug of unknown reliability, to gain information about its potentially beneficial effects. Mirman et al. (1993) find that a monopolist may set prices that deviate from the short term myopic optimal level, to collect information about the demand curve for its product. Rasmusen (2010) studies employee behavior if some tasks reveal their uncertain skills and other tasks do not. And Mueller and Scarsini (2002) show why risk averters should prefer uncertain over safe lotteries if learning is possible.

These studies show why people should choose uncertain alternatives so as to learn in a range of applications. However, they do not speak to the question of whether decision makers in the real world appreciate the benefits of learning

☆ We thank the editor, associate editor and two referees for helpful comments.
* Corresponding author. Fax: +1 617 384 9340.
E-mail addresses: s.t.trautmann@uvt.nl (S.T. Trautmann), richard_zeckhauser@harvard.edu (R.J. Zeckhauser).

0899-8256/$ – see front matter
http://dx.doi.org/10.1016/j.geb.2013.01.001

opportunities offered by alternatives with uncertain probabilities, and therefore choose these alternatives to collect valuable information. Indeed, the empirical literature suggests that people often shun uncertain alternatives, forgoing significant benefits. Muthukrishnan et al. (2009) show that consumers prefer products from well-established brands over cheaper products from less well-known brands, even when they consider them equally good in terms of quality. They find this effect for a large set of products, and find it most pronounced for people who avoid uncertain alternatives in a simple urn-choice task. In a study of patients with chronic conditions, Frank and Zeckhauser (2007) find that primary care doctors treating depression tend not to respond to changes in patients' symptoms when deciding whether to change medicines or dosages. Teodurescu and Erev (2011) discuss a wide range of situations – e.g., negotiations with uncertain mutual benefits, and the treatment of chronic depression – where people forgo the beneficial exploration of uncertain alternatives.

We hypothesize that the *neglect of learning opportunities* offered by uncertain probabilities applies broadly, and that it will be strongly manifest in choices made under laboratory conditions. Because the above empirical observations can be driven by factors unrelated to the neglect of learning benefits, we conducted a set of controlled experiments that aimed to eliminate influences such as trust, accountability, and liability, which often intrude on decisions outside the lab. All of our experiments involved compound lotteries with marbles in urns, or the equivalent. The urns were filled, as the subjects were informed, by random draws from larger known populations of marbles. Moreover, those larger populations were evenly divided between marbles of two colors. Finally, subjects were choosing for themselves. Thus, many extraneous factors that might intrude on the choice of a product or a medical treatment – such as information from elsewhere, past experience with an alternative, or agency concerns – were eliminated.

The previous literature studying decisions with uncertain probabilities mainly focuses on violations of expected utility (Savage, 1954) caused by the aversion to ambiguous and compound lotteries in one-shot situations (Ellsberg, 1961; Halevy, 2007; Kaivanto and Kroll, 2012; Spears, 2012). Learning benefits are not relevant to such choices.

Our analysis specifically allows for repeat choices and, therefore, for the opportunity for beneficial learning about uncertain alternatives when there are repeat draws from a stochastic process with unknown probabilities. We find that blindness to learning opportunities under uncertainty is deeply rooted. Alternatives with uncertain probabilities are typically shunned in favor of options whose probabilities are known. Our results suggest a range of explanations for subjects' violations of rational learning. In particular, four categories of behaviors lead people to violate optimal learning: underestimating learning, minimum learning, insensitive sticking, and fallacious switching. Participants making the first two errors underestimate learning opportunities although they get the direction of inference correct. Participants who insensitively stick or fallaciously switch, do worse: They make choices that are counter to what they have learned. These behaviors/errors relate to a set of requirements for rational behavior in potential learning tasks, which we introduce below.

Several papers provide surprising results showing insufficient appreciation of learning and counterintuitive results regarding the processing of feedback for equilibrium play in simple games (Merlo and Schotter, 1999; Rick and Weber, 2010). The current paper focuses differently than these studies: it makes anticipation of learning an inherent part of the decision problem. Consistent with the previous literature, however, we also find that there is little transferability across tasks. For example, forcing subjects to learn in one context does not lead them to choose learning in another. Our paper is closest in spirit to those of Charness and Levin (2005) and Charness et al. (2007), who show that violations of Bayesian learning are most common where the simple behavioral rule of reinforcement learning gives different predictions than Bayes' rule. Those papers consider decision problems where all probabilities are known with certainty. We address the basic learning paradigm when some probabilities are uncertain. We observe that people's lack of a clear understanding of learning under uncertainty leads them to behave in a non-Bayesian manner. In this sense our paper shows that the violations identified for known risks apply more broadly under uncertainty, and often prevent people from choosing a position that would provide them with beneficial learning opportunities.

This paper reports on a series of experiments designed to test if subjects anticipate and properly appreciate the learning opportunities offered by uncertainty, and to tease out the various types of violations of rational learning behavior. Each experiment offers an uncertain option that provides for formal learning in a laboratory decision task. Moreover, the benefits of such learning are subject to mathematical calculation. We did not expect any such calculations by our subjects, but they do permit calculations of expected value, and the loss of such value if the learning option is not chosen. As described below, we impose four lesser requirements for our subjects to behave in what might be called quasi-rational fashion, basically that they behave in a qualitative manner consistent with Bayes' Theorem, and the anticipation of learning. A third experiment requires a rough estimate of expected benefits, and thus imposes a fifth requirement of well-calibrated learning.

In Sections 1, 2, and 3, we present three incentivized lab experiments conducted with undergraduates at Tilburg University, the Netherlands. The experiments reveal both widespread shunning of options that provide learning benefits, and failures to learn in a rational fashion. Section 4 studies the limits of the learning violations and the transfer (or lack of transfer) of effective strategies across tasks. Section 5 discusses the implications of the results, and presents an extension of the current paradigm. Section 6 concludes.

1. Experiment 1: repeated urn[1] choices

1.1. Design

Ninety-nine undergraduate students participated in a laboratory experiment with real monetary incentives. To model a learning opportunity under uncertainty, we conducted a set of choice experiments utilizing two bags,[2] each with two colors of poker chips. In the baseline condition (BASE), we offered subjects a simple choice between a bet with known probabilities and an uncertain bet. Below, we often follow the terminology developed by Frank Knight (1921), where *risky* refers to bets with known probabilities, and *uncertain* refers to bets whose probabilities are not known. The risky choice was to bet on the color of a chip drawn from a bag known to contain 5 red and 5 black chips. If the subject guessed the color correctly, she won €10; otherwise she won nothing. The uncertain choice was to bet on the color of a chip drawn from a bag with 10 chips, either red or black, with the numbers of each color unknown. Again, a correct guess won the subject €10, otherwise she won nothing.

Both bags were assembled by the subjects themselves, who filled the bags with chips from a box with 50 red and 50 black chips. Each subject first filled the known bag, and then, wearing a blindfold, put 10 chips into the unknown-composition bag. If not stated otherwise, this procedure was followed in all urn-choice tasks in this paper. This transparent two-stage compound procedure was chosen to minimize any suspicion, and to emphasize the expected symmetry of colors in the uncertain option. It follows from the design that, in this baseline one-shot decision, the expected winning probability equaled 50% for both the known and the unknown-composition options, and the expected payoff from either option was thus €5. Note that given the equal numbers of the two colors in the box from which the unknown-composition bag was filled, and that draws were made without replacement, highly unbalanced distributions were quite unlikely.[3]

To introduce the potential for learning in this setting, we included the following repeated urn-choice task (REPEAT), in which the subjects chose between betting twice on the color of a chip drawn from the bag with known composition or making the same bets with the unknown-composition bag. The task had two components, A and B. For A, the subject picked either the known or the unknown bag. She would have to stick with this bag for the rest of the experiment. For B, she had to bet on two successive trials on the color of a chip drawn at random from her chosen bag. On each trial, she was paid according to her prediction, as described above. After the first trial and before making her prediction for the second trial, the chip was replaced in the bag. In the play on the first trial, she automatically observed the color of a chip drawn from her chosen bag. The subject then played the second trial, making a new prediction. She could stick with her original color or switch to the other color, for another €10 prize. Thus, she could potentially adjust her bet according to the information gained from the first draw. The repeated, two-trial structure of component B was explained to the subjects and illustrated schematically by a time line before component A of the game was decided, that is before they had to choose between the known-composition (risky) and unknown-composition (uncertain) bags.

In the repeated game, for the known-composition bag the probability of winning is 50% for each draw. This is also true for the first draw from the unknown-composition bag. For the second draw from the unknown-composition bag, however, a subject can increase her chances of winning because she learned something from the first draw. For instance, the drawing of a red chip on the first trial would be recognized by a rational subject as indicating that red is now more likely. Specifically, by predicting the color drawn from the unknown-composition bag the first time as her second prediction, the subject will win 54.5% of the time on the second draw.

While in the BASE experiment both options are equally good from a Bayesian perspective, in the REPEAT game the expected value of the uncertain option is €10.45 versus €10 for the known option. A rational person might still pick the known option. She might make a trade-off between learning benefits offered by the unknown-composition bag, and other motives – e.g., pure uncertainty aversion or simply a preference not to make the cognitive effort to worry about learning – that led her to shy away from uncertainty (Ellsberg, 1961; Halevy, 2007; Kaivanto and Kroll, 2012; Rosenboim et al., 2013; Spears, 2012). Thus, she may not choose the uncertain alternative even if she correctly understands learning. But if this is the case, increasing the relative benefit of the learning opportunity would make it more likely that the unknown-composition bag would be chosen.

In contrast, if a subject does not recognize the benefits of learning under uncertainty, increasing the relative benefit of learning opportunities would not affect her choice between the known and unknown bag. To discriminate between these two explanations, we employed a second two-trial treatment to see whether increased learning benefits made a difference. In treatment REPEAT4, the known bag contained exactly two red and two black chips, while the unknown-composition bag contained four chips, either red or black but in an unknown proportion. Otherwise the choice task was identical to the 10-chip, repeated, urn-choice task described above. Urns with fewer chips provide more significant learning, as going to the extreme case of a 1-chip urn makes clear. For the 4-chip case, the expected payoff from the known option is again €10,

[1] We use the terms "urn" – widely employed in the decision analysis literature – and "bag" interchangeably in this paper.

[2] In the experiments we used bags instead of urns. We will use the terms bag and urn interchangeably in the general description of choice situations. In some experiments, marbles were used instead of poker chips.

[3] Note that the procedure of filling the urn is irrelevant to the *qualitative* learning aspects, which we formalize in Requirement 3 (Bayesian Tendency) below. As soon as the urn has been filled, sampling is informative about the actual, uncertain distribution of chips in the urn.

Table 1
Choice of uncertain bet.

	BASE	REPEAT10	REPEAT4	REPEAT (combined)
N	31	33	35	68
Uncertain option chosen	7 (23%)	13 (39%)	10 (29%)	23 (34%)
Binomial test, two-sided	$p = 0.003$	$p = 0.296$	$p = 0.017$	$p = 0.010$

while for the uncertain option the expected payoff now equals €11.20, a 12% increase in expected payoff from correctly choosing on the second trial.[4]

1.2. Requirements for quasi-rationality

This experiment well illustrates the requirements for individuals to respond effectively to our questions, and more generally to deal appropriately with uncertain alternatives. It is important to note that none of these requirements entails quantitative calculations. We identify four requirements for rational decision making here, and one in Section 3, and identify their prime accomplishments.[5]

Requirement 1. *Compound Lotteries*. Individuals reduce compound lotteries.
Requirement 2. *Independence*. Individuals recognize that successive draws with replacement are independent.
Requirement 3. *Bayesian Tendency*. Individuals draw qualitative inferences about the composition of an urn from the first draw and make choices that are consistent with Bayes' rule.
Requirement 4. *Forward Looking*. Individuals look ahead to the inferences they will make after observing a draw from the unknown urn.

Requirement 1 frees individuals from various violations of expected utility in static-choice settings.[6] Requirement 2 avoids the gambler's fallacy: assuming that a color that had not come up the first time was more likely to come up on the second trial, when sampling with replacement. Requirement 3 assures that individuals make choices consistent with clear learning, as would come from a single draw of a marble. We shall also discuss Requirement 3′, its subcomponent, which is that when there is clearly more learning individuals recognize that. Requirement 4 implies that individuals will make the correct decision initially when they have to choose between the known and the unknown-composition urn. Note, Requirement 3 overlaps separately with Requirement 1 and Requirement 2. Hereafter, requirements will usually be abbreviated as R1, R2, R3, and R4.

1.3. Results

Table 1 shows the results. In the baseline one-shot treatment, we replicated the pattern found in previous studies where options with known probabilities are preferred to options with uncertain probabilities that offer the same expected value. Only 23% of the subjects chose the uncertain option.[7] No preference for the (strictly better) uncertain option was observed in the two-trial treatments offering learning. The uncertain option received less than 50% of the play in both the REPEAT10 and REPEAT4 treatments. Putting the results from the two repeat treatments together, the choice of the unknown-composition urn is strongly significantly less than 1/2, the expected fraction if subjects merely chose at random ($p = 0.010$, binomial test). This pattern represents a severe violation of the joint hypothesis that R3, Bayesian Tendency, and R4, Forward Looking, will be met.[8] Thus, subjects who either mislearned or neglected learning or those who were myopic might choose the known urn.

The REPEAT4 bag provides greater learning than the REPEAT10 bag, since one learns about a greater percentage of the chips. (The gain in expected value is 2 2/3 times as much, 1.20 as opposed to 0.45.) However, REPEAT4 was selected less often, 29% versus 39%, contrary to R3′, possibly because it is not intuitively obvious to some that it offers more learning

[4] For the simple baseline game with no learning opportunities, it has been shown that the size of the urn (that is, the number of chips) does not affect behavior in one-shot choices between known and uncertain options (Pulford and Colman, 2008).

[5] We thank a referee for suggesting that we set forth requirements for rational decision making and test whether they are met singly and jointly, and for clarifying our discussion of these requirements.

[6] Decision makers who fall prey to the Ellsberg Paradox violate Requirement 1. See Segal (1990) for a detailed discussion of this requirement. Clearly, Requirement 1 is not sufficient to prevent all violations of expected utility in static settings.

[7] This pattern is perfectly consistent with full rationality, since there was no benefit in BASE for choosing the uncertain option, but it does suggest a tendency to shun uncertainty.

[8] Were R3 and R4 satisfied, the observed pattern would not be possible. A failure of R1 or R2 could be the source of R3's failure. However, if R1's failure were the sole explanation, given R3′ (when there is more learning individuals recognize that), REPEAT4 would probably have the uncertain urn chosen more than REPEAT10.

Table 2
Learning behavior in repeated-bet game.

	Option chosen	Stayed with color predicted in stage 1		Compatible with learning
		First color right	First color wrong	
10-chip repeated	Uncertain	5/7 (71%)	3/6 (50%)	8/13 (61%)
	Known	8/11 (73%)	9/9 (100%)	8/20 (40%)[AIL]
4-chip repeated	Uncertain	3/5 (60%)	2/5 (40%)	6/10 (60%)
	Known	10/14 (71%)	10/11 (90%)	11/25 (44%)[AIL]

Note: The AIL superscript indicates "as-if learning" for the known-risk option: a choice that would have been compatible with learning had the subject played the uncertain option.

than the 10-chip bag.[9] However, the difference between these two treatments was not statistically significant ($p = 0.444$, Fisher's exact test, two-sided).

Table 2 provides strong additional evidence regarding the neglect of the potential learning benefits from uncertainty. For the repeated-trial treatments, this table shows the number of subjects who stayed with their first-trial prediction after either a successful or an unsuccessful prediction. Of the subjects who chose uncertainty, almost 40% behaved directly contrary to learning, and in accord with the gambler's fallacy, a direct violation of R2. They fall into two groups. Of those subjects who guessed correctly on the first trial, many (29% in REPEAT10 and 40% in REPEAT4) switched colors, although the initial success indicated their initial color was more likely on the second trial. We label these individuals fallacious switchers. Of those guessing wrong on the first trial, many (50% for REPEAT10 and 40% for REPEAT4) stuck with their initial prediction, despite information suggesting the other color was more likely. We call these people insensitive stickers. Both fallacious switchers and insensitive stickers behave in accord with the gambler's fallacy, an indirect violation of R2, Independence. In doing so, they also violated R3, Bayesian Tendency.

Table 2 also shows that subjects who chose the known option sometimes switched after an initial success, but almost never switched after an initial failure. That is, they deviated from purely random choice. Though this entailed no cost, given the known and balanced composition of the urn, it is the behavior one prone to the gambler's fallacy would follow.

To our knowledge, there is no experimental evidence on repeated choices using a chance device between known risks and uncertain risks that offer learning opportunities. Yet, we believe that the potential to learn from uncertain situations is the norm in the world, albeit not with such artificial situations as drawing a ball from an urn. One frequently has a repeat choice of dry cleaners, driving routes, or restaurants, or strategies to employ when dealing with employees or friends. In each case, there will be available options where you are highly uncertain about outcomes, but it is logical to believe that what happens on the first trial will correlate positively with the outcome on a second trial, implying the opportunity for learning. The driving route question is as follows: "You commute to work every day. There are two possible routes. Route A, your usual route, takes 30 minutes on average with a standard deviation of 3 minutes. You have only tried Route B three times. It has taken 28, 31, and 34 minutes. You expect to be commuting for many years. Should you stick with Route A, or try Route B a few more times?" We have asked this question in many venues. Few individuals recognize the two-armed bandit feature of the problem, and the advantage of gathering information on Route B.[10]

We should note that some past experiments involving no learning possibility have shown a tilt toward uncertainty in repeated settings. Liu and Colman (2009) conducted experiments with repeated bets, each time with a newly assembled urn (hence no learning), and no switching opportunities. If the unknown-composition urn offers a larger prize and 100 repetitions are made, subjects prefer the unknown-composition urn, though they preferred the known urn on one trial. These authors argue that, through repetition, the uncertainty about the probabilities is reduced to a 50% chance as in the known-risk lottery, but with a higher prize.

Rode et al. (1999) studied the case in which at least x red balls have to be drawn in n trials. Indeed, they found that people correctly prefer uncertainty if x is large relative to expectation, implying that the known urn would give a low chance of achieving x successes, so that the unknown-composition urn would offer a higher expected payoff. We will present some data on this paradigm in Section 5.

Our results on switching strategies replicate basic patterns observed by Charness and Levin (2005) in decisions that involve only known probabilities. Many subjects violate optimal switching strategies, and these violations do not diminish with the stronger learning potential in the 4-chip case (see their result 3, p. 1305). In our setting, such learning violations indicate a much stronger bias against the favorable urn. 60 to 70 percent of our subjects chose the less favorable urn. In their study, a pure known-risk setting, only 20 to 30 percent of the subjects chose the less favorable urn.

2. Experiment 2: making learning opportunities salient

In Experiment 1, we identified an insensitivity to learning, reinforced by a strong type of gambler's fallacy. (Subjects choosing the known-risk option, thus learning nothing, had no opportunity to fall into this costly error.) Furthermore, the

[9] It violates 3′ in the sense that the tendency for REPEAT4 should be stronger. Of course, if all individuals were perfect Bayesians, 100% of both groups would choose the uncertain option.

[10] Indeed, this advantage would persist even if Route B's average time to date were slightly greater.

Table 3
Predicted composition and urn choices after forced sample.

	Part II urn choice		
	Known	Uncertain	All
Urn choice	30 (64%)	17 (36%)	47
Part I: violate learning in prediction of bag composition	7 (23%)	5 (29%)[a]	12 (26%)
Part II: violate learning in color prediction	–	4 (24%)[a]	
Total incidence of learning mistakes	7 (23%)	7 (41%)	14 (30%)

[a] These two groups had two members in common.

learning behavior emerged endogenously from observations of outcomes in the uncertain option, similar to the ways learning under uncertainty emerges in real-world decisions. Possibly the participants merely overlooked the learning opportunity. While such obliviousness might apply in many real-world problems, we predicted that if the possibility for learning were made more salient, people might correctly perceive strong benefits from pursuing it. They would therefore predominantly choose the alternative with uncertain probabilities. To test this, we conducted an experiment in which people were forced to observe a sample before making their choices between known and uncertain options, and where we could identify whether those choosing the risky option also learned.[11]

2.1. Design

Forty-seven undergraduate students participated in an experiment that built on the BASE condition described in the previous section. The experiment consisted of two parts. Part II was identical to the 10-chip bag, one-shot, two-color bet for a prize of €10 in BASE: choose between the known and the unknown bag, and bet on the color of a chip once. In Part I, subjects had to draw one chip from the known bag and one chip from the unknown-composition bag, always with replacement. They noted the colors sampled, and then had to predict the contents of the two bags, that is, whether they expected more red chips or more black chips, or an equal number of red and black chips in each one. If the predicted compositions of both bags were correct, they would win €10. It was made clear to the subjects that exactly the same bags with the same contents that they had sampled in Part I would also be used in the BASE task in Part II where they had to choose one bag and bet on a color. At the end of the experiment, one of the two parts would be chosen by coin toss for real payment.

For the unknown-composition bag, sampling a red chip implied that the bag was more likely predominantly red than either predominantly black or equally distributed. For the known bag, the question was trivial because subjects knew that it contained equal numbers of red and black chips. The question was included for reasons of symmetry, and to check basic understanding of the procedure. Indeed, all subjects correctly indicated the equal distribution in the known bag.

Sampling from the unknown-composition bag in Part I of the experiment forced the subjects to observe information about the distribution of colors, and allowed them to upgrade their expected likelihood of winning the prize for the Part II urn-choice task. After sampling red, the probability of winning a bet on red increased to 54.5%. The expected value of the ambiguous option was €5.45 versus €5.00 for the risky option, an increase of 9%. Note that the benefit from learning was larger in this experiment than in the REPEAT task in Experiment 1. In the repeated game, the subjects would still have to make the first bet without any information; hence, the proportional effect on total expected earnings would be smaller. Because all the subjects had to sample the unknown-composition bag and make a prediction regarding its contents, we could also observe learning errors for those choosing the known option.

2.2. Results

Table 3 shows the results. We found a similar level of unknown-composition urn choices as in the previous experiments, with 36% choosing uncertain ($p = 0.079$, binomial test, two-sided). Overall, in the prediction of the contents of the unknown-composition urn, 26% violated learning and did not predict a majority of the color they drew, and thus violated learning and R3, Bayesian Tendency. The incidence of this violation was similar among known-risk and uncertain-risk choosers. For unknown-composition urn choosers, we also observed whether they violated learning, and thus R3, by betting adversely on the color not drawn in the sampling draw. We found that 24% committed such errors and that this group did not completely overlap with the group of people who failed the predominant color prediction in Part I of this experiment.

35 subjects – those represented in the bottom two rows of the table – stated that they believed that the color they picked in Part I in the bag with uncertain composition was in the majority. Of this group, 23 (roughly two thirds) then chose the known-composition bag in Part II, apparently in violation of R3, Bayesian Tendency. We asked these subjects why they had not chosen the uncertain option, and then bet on the color they believed predominated in the bag. Many argued that they had not perceived the sample as strong evidence for the color drawn, and some indicated that they had

[11] Charness and Levin (2005) use a similar manipulation where subjects are forced to make a specific choice in the first of a series of choices but are allowed to switch after the first draw.

predicted the majority composition more or less randomly. Since they had seen no clear evidence for either color, they had preferred the known option in the Part II bet.[12] However, as the data show, it was still common for the subjects to predict the contents according to their samples. Those who correctly predicted the composition but then shunned the uncertain option might basically have had the right intuition. If so, they notably underestimated the value of the sample.

Another group of subjects fall into the class we call minimum learners. They also announced the learning-compatible contents, but then chose the known option. When asked about their decision, these subjects suggested a maximin way of thinking about probabilities: drawing a sample of a red chip was counted as evidence of one red chip versus no evidence for black at all. Such reasoning implied a correct majority prediction because there was more evidence for the red, and at the same time a preference for the known bet, because there were assuredly 5 red chips in the known option versus assurance of only 1 red chip in the uncertain option.[13] These individuals violated R3.

3. Experiment 3: putting a price tag on benefits from learning

Experiment 2 replicated the strong violations of learning that we found in Experiment 1. It suggested in addition that even those who do not commit such strong violations may still miss the clear learning potential offered by uncertain probabilities. To study how well people are calibrated when learning under uncertainty, we designed an experiment in which the subjects would not choose between known and uncertain bets. Rather, they would only make bets on uncertain options, some that offer learning and some that do not.

To deal effectively with this experiment, individuals would have to be able to intuitively approximate the magnitude of the learning opportunity, a more challenging requirement than the four listed above. Thus, we had:

Requirement 5. *Well-Calibrated Learning*. Individuals will be able to roughly approximate the degree of learning from a single draw from an unknown urn.

3.1. Design

Forty-three subjects participated in an experiment in which they had to predict the color of a marble drawn from an urn with 4 marbles, either red or black in an uncertain proportion.[14] Without any sample, the probability of winning the prize in this bet equals 50%. With a sample of one marble with replacement, predicting the color sampled increases the chance of winning to 61.73% and the expected value of the gamble by 23.47%. The percentage benefit from learning is larger in this experiment than in the REPEAT4 condition with 4 chips in Experiment 1. In the repeated game, the subjects would still have to make the first bet without any information; therefore, the effect on the total expected earnings is smaller.

To measure the strength of the learning opportunities perceived by the subjects, we elicited a prize-equivalent for sampling as follows. Subjects could either bet on a color drawn from the uncertain 4-chip bag, without any sample, for a winning prize of €20 for a correct prediction, or bet on a color drawn from this bag, after sampling one chip with replacement, for a winning prize of €x. There existed some $x < 20$ such that the subject would be indifferent about either directly betting for a prize of €20 or betting after learning something about the distribution of colors for the lower prize of €x. We called this indifference value the *lowest-acceptable prize* (LAP) of the sampling opportunity.

We elicited the LAP using a Becker–DeGroot–Marschak (1963) (BDM) mechanism. We placed 39 slips of paper in a bag with prize offers between €0.50 and €19.50 in equal steps of €0.50, and announced this composition to subjects. Subjects wrote down their LAP. If the randomly selected offered prize "y" was equal to or larger than the specified LAP, the subject would make a bet after sampling once for a prize of €y. If the offered prize "y" was smaller than the LAP, the subject would directly predict the color of a chip, without sampling, for a prize of €20. After writing down their LAP, subjects also had to specify the color they wanted to predict in case they played without a sample for a prize of €20, and in case they played with sampling for a price of €y. In the latter case, they had to specify two predictions, conditional on the color sampled. That is, we elicited full betting strategies for all contingencies.

Under expected payoff maximization, the optimal LAP was €16.20 because winning €16.20 with a 61.73% chance offers an expected value equal to winning €20 with a chance of 50%, in the case of no sampling opportunity.[15] The optimal strategy, obviously, involves betting on the color sampled.

3.2. Results

Based on the optimal LAP of €16.20, we defined people as well-calibrated learners, and hence meeting R5, if they specified a LAP between 14 and 18 inclusive. That is a range of two full Euros below and above the optimal value, and

[12] As with Experiment 1, a violation of R1 or R2 could lead to R3's violation. Subjects' explanations of their choices suggested other sources.

[13] The single draw of a red chip conveys considerable information about other possible combinations of bag contents. For example, the odds of 7 red chips and 3 black versus 3 red and 7 black, which were originally even, have now shifted to 7 to 3.

[14] Experiment 3 was conducted with marbles rather than chips. The marbles were drawn from a box containing 25 marbles of each color, instead of 50 of each, as in the previous experiments. Calculations account for the size of the box.

[15] Risk aversion should lower a rational Bayesian subject's LAP below 16.2. A 50% chance to win 20 Euros is a mean-preserving spread of a 61.73% chance to win 16.20 Euros, hence is more risky.

Table 4
Valuation of learning opportunity.

		Well calibrated $14 \leqslant \text{LAP} \leqslant 18$	Too hesitant to take sample (LAP > 18)	Too eager to take sample (LAP < 14)	All
All	Distribution by LAP	20 (47%)	8 (19%)	15 (35%)	43
	Avg. LAP	€15.73	€19.38	€9.07	€14.08
Correct strategy	Distribution by LAP	16 (57%)	5 (18%)	7 (25%)	28
	Avg. LAP	€15.91	€19.20	€9.71	€14.95
Insensitive stickers	Distribution by LAP	2 (25%)	2 (25%)	4 (50%)	8
	Avg. LAP	€15.00	€20.00	€9.50	€13.50
Fallacious switchers	Distribution by LAP	2 (29%)	1 (14%)	4 (57%)	7
	Avg. LAP	€15.00	€19.00	€7.5	€11.29

Note: Percentages refer to the distribution within the row.

it includes the prominent amount of €15.[16] Thus, we applied a conservative criterion for poorly approximated learning. Table 4 shows the results.

The first row shows the distribution of LAPs. Roughly half of the subjects were well calibrated (LAP between 14 and 18); about one fifth were too hesitant to sample (LAP too high); and 35% were too eager (LAP too low). Those too hesitant and those too eager both violated R5.

The results in the too-eager group look surprising, because they suggest a great value was placed on learning, whereas Experiment 1 showed that individuals even neglected learning that was costless. The explanation may fall in two parts. First, the too-eager individuals were explicitly required to value learning benefits. Second, they were terrible in doing so. Indeed, their average LAP was below the expected value of the gamble assuming the sample led to a guess that was always correct. This weak performance is reinforced because a remarkable 8 out of the 15 subjects who overpay for the draw then fall prey to the gambler's fallacy. That is, they violate R2 and then choose the wrong color on the basis of the draw they receive. Thus, they overpay for learning, and then neglect it.

Note that the LAP implies a comparison between an uncertain option with no sample and an uncertain option with a sample of one draw. In contrast, the previous experiments compared an uncertain option with a sample of known risk with 5 chances out of 10. Subjects who held some maximin view might feel that they learned very little from one sample, but would find that far superior to no identifiable knowledge (unknown-composition urn with no sample). Thus, they strongly preferred the one-draw alternative and were too eager to reduce the uncertainty. In interpreting these results with the BDM mechanism, we issue a cautionary note. BDM may be difficult for participants to understand, and the "overpayment" for learning may have come because individuals did not grasp how they should respond.[17]

The remaining rows of Table 4 distinguish subjects according to their betting strategies. Individuals who employed the correct learning strategy had a mean LAP about €15. They were well calibrated 57% of the time. There were two groups who specified strategies contrary to the correct inference from the potential outcomes of the sample. 8 of the 43 (19%) subjects specified a color that was independent of the outcome of the sample (insensitive stickers). 7 subjects (16%) specified a bet against the color sampled (fallacious switchers). Interestingly, the latter group had lower average LAP, €11.29, than those following correct strategy (€14.94; $t(33) = 2.203$, $p = 0.035$). Thus, they picked LAPs as if they were learning a lot, but then picked the wrong color. This group fell prey to the gambler's fallacy. Such choices are well documented for choices with known probabilities, as say with red and black on a roulette wheel. Here they are more disturbing, because they are going contrary to learning. Roughly two thirds (28 of 43) of subjects bet correctly. And 37% of subjects were both well calibrated and followed optimal betting.

4. Experiment 4: the limits of learning neglect

The previous experiments demonstrate various failures to anticipate or utilize learning under uncertainty when that uncertainty cannot be completely resolved. To examine the limiting conditions of this phenomenon, we considered two variations that we predicted would reduce the incidence of neglected learning benefits. First, we hypothesized that learning opportunities that eliminated all uncertainty would be taken by the subjects. Second, we predicted that experience in a learning task that revealed the general principle of learning under uncertainty would transfer to decisions in a task where learning was less obvious to people, the types of situations considered above.

4.1. Design

Thirty-two subjects participated in an experiment that had two parts. Each part involved monetary incentives and was presented as a separate experiment to the subjects. At the end of the experiment, one part was randomly selected by a

[16] Most people specified full Euro amounts.

[17] On average, our participants were willing to accept 14.08 for the game with learning, as opposed to 16.20, the actuarially fair value. They overpaid.

Table 5
The limits of learning neglect.

Part 1:	Known (white: 1 red and 1 black)		Unknown (blue: either two red or two black)		All	
	8 (25%)		24 (75%)		32	
Part 2:	Known	Unknown	Known	Unknown	Known	Unknown
	7 (88%)*	1 (12%)*	15 (63%)*	9 (37%)*	22 (69%)*	10 (31%)*
Compatible with learning (within Part 2)[a]	1 (14%)**	1 (100%)**	4 (26%)**	9 (100%)**	5 (23%)**	10 (100%)**

Notes: Part 1: Repeated bet with possible resolution of uncertainty. Part 2: Repeated 4-chip urn-choice task REPEAT4.
[a] "As-if learning" for known-risk option: Indicates a choice pattern that would have been compatible with learning had the subject played the unknown-risk option.
* Probabilities conditional on Part 1 choice.
** Probabilities conditional on Part 1 and Part 2 choices.

coin flip for real payment. The second part of the experiment was identical to the REPEAT4 repeated urn-choice experiment presented in Section 1. The first part of the experiment involved the following choice situation, modeled on the choice task in Charness and Levin (2005). Subjects were presented with one white bag and two indistinguishable blue bags. The white bag contained exactly one red and one black marble. One of the blue bags contained exactly two red marbles; the other blue bag contained exactly two black marbles. The subjects knew the possible contents of the two blue bags, but did not know which one contained only red marbles and which contained only black. Before learning about the decision problem, subjects chose one of the blue bags for the experiment, along with the white bag. The other blue bag was removed.

The decision problem was similar to the learning setting in Section 1. The subjects had to choose either the white bag with the known and equal distribution of red and black, or an uncertain blue bag containing only either two red marbles or two black marbles. The bag would be used to make a repeated bet on the color of a marble drawn, with replacement. Specifically, each subject first chose a bag, predicted a color, and then drew a marble; and the subject won €5 if the prediction was correct, and nothing otherwise. The marble was replaced in the bag, and the subject then again chose a bag, predicted a color, and drew a marble, winning another €5 if the prediction was correct. In this repeated-betting situation, the chance of winning the prize equaled 50% for the known-risk white bag in both the first and the second drawings. For a blue bag, the first drawing also offered a 50% chance of winning, since the bag had been randomly selected with an equal chance of containing either two red or two black marbles. After betting on a color drawn from the blue bag in the first drawing, however, the subject learned the marble's color with certainty. That is, after the uncertainty of the first draw, the blue bag offered a certain gain of €5 in the second draw.

After the first part of the experiment was finished, the bags were set aside; then Part 2, the repeated urn choice REPEAT4, was conducted exactly as described in Section 1. The goal was to see if the definitive learning situation with the blue bags would help subjects recognize the potential for valuable but imperfect learning with the unknown-composition bag.

4.2. Results

Table 5 shows the results. In the Part 1 repeated-betting task with complete resolution of uncertainty for the uncertain blue bag in the second drawing, 75% of the subjects chose the uncertain option ($p = 0.007$, binomial test, two-sided). That is far superior to the previous experiments; the majority correctly understood the learning opportunity available for blue. Given the simplicity of the task, it is somewhat surprising that at least 25% of subjects did not see the benefit of the uncertain probability.[18] Their choice represents a stark violation of R3 or R4 or both.

The presence of this group of failed learners is consistent with the incidence of such learning deficits in our previous experiments, and also with the numbers observed in Charness and Levin's (2005) baseline condition, the basis for the design of the Part 1 decision. The results show that the potential for complete resolution of uncertainty can reduce but hardly eliminate learning neglect. For subjects who intuited the basic advantage but underestimated the magnitude of learning, and for those who applied maximin learning, this task's structure immediately revealed that uncertain probabilities offer learning opportunities.

While 75% of subjects acted as if they identified the learning opportunity in Part 1, the experience with Part 2 is disappointing. Only 31% of subjects chose the uncertain option in the repeated 4-chip urn-choice task REPEAT4 ($p = 0.050$, binomial test, two-sided), despite their exposure to Part 1. Indeed, the percentage of people choosing the uncertain option differed little from that in Experiment 1. Clearly, the Part 1 experience failed to allow or alert subjects to understand the basic concept that securing information on an uncertain probability provides information. Most subjects who chose the known-risk option in Part 1 also chose the known risk in Part 2, as would be expected.

Our major interest is in the 24 individuals who chose the uncertain option in Part 1. They met R4, and absent further knowledge presumably understood the Forward Looking strategy. Further investigation disappoints. A surprising 63% of them chose the known-composition option in Part 2. Thus, a majority of the subjects who appeared to get the right intuition in

[18] 25% is a lower bound, since some who chose the blue bag may have done so at random.

Part 1, with certain learning, did not successfully transfer their insights regarding learning opportunities to the partial learning task of Part 2. Their choices surely refuted the joint hypothesis R3 and R4. However, a failure on any one of R1, R2, R3, or R4, by itself, could have led to their Part 2 choice.

Although the Part 1 experience was not sufficient to make most subjects recognize the benefits of uncertainty, it had some effects on behavior in the Part 2 decision problem REPEAT4. First, of those subjects who chose the unknown bag in Part 2, all made choices according to optimal learning, switching to/staying with the color drawn in the first trial (see bottom row of Table 5). In Experiment 1 with no previous exposure to a statistical learning task, only 60% of subjects chose correctly in this fashion (Table 2).[19] On the other hand, the Part 1 exposure seems to increase the incidence of gambler's fallacy behavior for those who chose the known-composition bag in Part 2. These risky-choosers learned nothing in Part 2 between the two trials. However, the "as-if learning" entry in the bottom row of Table 5 describes their switching behavior between trials and thus serves as a benchmark for the behavior of the subjects who chose the unknown bag in Part 2. Table 5 shows that only 23% of the risky choosers ($p = 0.017$, binomial test, two-sided) switch as would be optimal had they chosen the unknown-composition bag. Because the index is below 50%, it shows that subjects switch after success and stay after failures, as is predicted by the gambler's fallacy. Thus, while for some people the Part 1 task helps them to identify the optimal learning strategy, and to benefit from the uncertainty about probability, for a larger group of people it seems that Part 1 potentially added to their confusion about optimal learning, leading to widespread gambler's fallacy behavior for the risky bag.

5. Discussion

In the real world, uncertain probabilities of success are often accompanied by a learning opportunity because they will have future choices whose outcomes will be correlated. If an innovative fast food restaurant proves successful in one neighborhood, it is much more likely to be successful in a second similar neighborhood. An employee who shows himself capable in one context is much more likely to be capable in another similar context. Real-world decision makers do not seem to grasp the conceptual underpinnings of such learning opportunities. They often do not seem to anticipate the benefits of learning, and therefore do not choose alternatives with uncertain probabilities. They are likely to choose option A, which has a 60% chance of succeeding over option B, which has a one half chance of being at 90% but one half of being at 20%. This makes sense when a choice is made once, since the probability of success is 60% rather than 55%. But the decision maker should "pay for learning" and take the uncertain alternative given these values when there are two or more choices to be made. Failing to experiment in this way, they will not observe successful outcomes from uncertain alternatives. Thus, they will have no chance to benefit from learning.

We believe that most decision makers would properly interpret the observation of 4 successes and 1 failure as evidence of a high success probability if dealing with an uncertain 10-chip urn. However, our results strongly suggest that most people will not obtain helpful samples because, due to neglect or misconception of the mechanics of learning, they will shun uncertainty, settling on known alternatives that offer few or no learning opportunities. Examples where learning is not possible without an initial commitment to uncertain alternatives include scalable investments, new varieties of seed, or new medical treatments.

From an evolutionary perspective, if the benefits from learning are large, why would learning avoidance persist? This is a profound question. Fortunately, some research points to a potential direction for the search for an answer. Psychological findings suggest that negative experiences are crucial to learning, while good experiences exert virtually no pedagogic power (Baumeister et al., 2001). This appears related to the finding that in individual decision situations, losses often weigh 2 to 3 times as much as gains (Tversky and Kahneman, 1992; Abdellaoui et al., 2007, Table 1, p. 1662). In the current setting, uncertain options would need to be sampled repeatedly in order to obtain a sufficient sample with few negative outcomes to determine whether to switch from the status quo. People require too much positive evidence before shifting to uncertain options.[20]

Other considerations, such as regret and blame avoidance, may also contribute to shunning uncertain alternatives. One does not know what returns would have come from an uncertain alternative. This reduces regret from not having chosen it. Blame from others also plays an important role. In principal-agent relationships, bad outcomes often lead to criticism, and possibly legal consequences because of responsibility and accountability. Therefore, agents, such as financial advisors or medical practitioners, may weight bad relative to good payoffs more than do their principals (Eriksen and Kvaloy, 2010a, 2010b). Most people, for that reason, have had many fewer positive learning experiences with unknown alternatives than rational decision theory would prescribe.

Our results may complement our understanding of herding behavior and behavioral contagion in financial markets (Hirshleifer and Teoh, 2009). In a recent experimental study, Goeree and Yariv (2006) brought subjects into a situation with uncertain probabilities similar to the ones studied in our paper. They then let them choose between an informative private signal, and an uninformative social signal. Specifically, subjects had to predict the contents of a jar that was filled

[19] 10 out of 10 subjects in Experiment 4 behaved compatible with learning, while 14 out of 23 behaved compatible with learning in Experiment 1 ($p = 0.032$, Fisher's exact test, two-sided).

[20] We recognize this is not a complete explanation of the outcome from evolution. The question then becomes why we need negative outcomes to learn.

with balls that were either predominantly red (7 red and 3 blue) or predominantly blue (7 blue and 3 red). The prior probability of either distribution was 50%. Subjects could choose one of two pieces of information: (1) they could sample once with replacement from the jar before making their guess (informative statistical signal); or (2) they could choose to be told the predictions of 3 people who before them had randomly guessed the distribution without any statistical signal (uninformative social signal.). Across different conditions Goeree and Yariv find that between 34% and 51% of their subjects choose the uninformative social signal. They conclude that an intrinsic taste for conformity can explain their result.

If people correctly understood the benefits of learning, Goeree and Yariv's result would imply a strong preference for conformity. A study by Corazinni and Greiner (2007) using a simple risky choice paradigm instead of learning, questions whether such strong preferences for conformity exist. Our results suggest that even a weak preference for conformity may be enough in the Goeree and Yariv learning paradigm, thus reconciling their results with those of Corazinni and Greiner. We showed that most people have little concept of learning in situations where information is gained but uncertainty is not definitively resolved. These learning violators will not perceive the statistical sample as a valuable option, implying that mild curiosity or conformity would be enough to induce them to copy other people's uninformed choice.

An interesting extension of the current paradigm concerns situations where uncertain alternatives are encountered repeatedly, and a success is needed on each trial to guarantee an overall success. For example, a proposal in an organization might need a signoff from four independent divisions before it can go forward, with an equal probability of success in each division. The traditional format proposal has a 70% chance of approval from each division, implying an overall success rate of slightly less than 25%. The new format being contemplated has an uncertain chance of success within each division. It could be 90% or it could be 40%. Each is one half likely. Its overall success probability is $\frac{1}{2}(0.9^4 + 0.4^4) = 33.6\%$. Though the percentage for success at each division is lower, namely $(90 + 40)/2 = 65\%$, the overall probability of success has increased. This setting is quite different from the above learning paradigm. The positive correlation in the probability of success across the four divisions is the source of the increase.

Given the results of the current paper, we predicted that people would have difficulties understanding the benefits of uncertainty in this quite different repeated setting as well. We conducted an experiment using the bags filled according to Experiment 1's REPEAT10 conditions. The subjects had to pick the known or unknown bag, and then pick one color that would apply for two draws with replacement. If both draws were of their color they won 10 Euros, and otherwise nothing. The chances of winning were 25% with the known urn, but 27.2% with the unknown urn. Only 12 out 32 subjects chose the superior unknown-composition urn. The potential beneficial effects of uncertain probabilities when a series of successes is needed were missed. Uncertainty was shunned quite apart from the learning paradigm studied in this paper.

6. Conclusion

Whether in financial, medical, or other decisions, learning opportunities in which outcomes and probabilities are uncertain offer large expected gains over known risks. Many paths produce erroneous thinking about learning; the most prominent of them simply does not see the possibility. Indeed, an extreme learning experience – full resolution of a probability – in a similar setting proved to be an insufficient spur. The broad finding from diverse experiments is that individuals shun uncertainty and fail to recognize the benefits of learning that it offers. They fail to meet the basic requirements of rational decision making.

Supplementary material

The online version of this article contains additional supplementary material.
Please visit http://dx.doi.org/10.1016/j.geb.2013.01.001.

References

Abdellaoui, M., Bleichrodt, H., Paraschiv, C., 2007. Loss aversion under prospect theory: A parameter-free measurement. Manage. Sci. 53, 1659–1674.
Baumeister, R.F., Bratslavsky, E., Finkenauer, C., Vohs, K.D., 2001. Bad is stronger than good. Rev. Gen. Psych. 5, 323–370.
Becker, G.M., de Groot, M.H., Marschak, J., 1963. Stochastic models of choice behavior. Behavioral Sci. 8, 41–55.
Charness, G., Levin, D., 2005. When optimal choices feel wrong: A laboratory study of Bayesian updating, complexity, and affect. Amer. Econ. Rev. 95, 1300–1309.
Charness, G., Karni, E., Levin, D., 2007. Individual and group decision making under risk: An experimental study of Bayesian updating and violations of first-order stochastic dominance. J. Risk Uncertainty 35, 129–148.
Corazinni, L., Greiner, B., 2007. Herding, social preferences, and (non-)conformity. Econ. Letters 97, 74–80.
Ellsberg, D., 1961. Risk, ambiguity, and the savage axioms. Quart. J. Econ. 7, 643–669.
Eriksen, K.W., Kvaloy, O., 2010a. Do financial advisors exhibit myopic loss aversion? Financial Markets Portfol. Manage. 24, 159–170.
Eriksen, K.W., Kvaloy, O., 2010b. Myopic investment management. Rev. Finance 14, 521–542.
Frank, R.G., Zeckhauser, R.J., 2007. Custom-made versus ready-to-wear treatments: Behavioral propensities in physicians' choices. J. Health Econ. 26, 1101–1127.
Goeree, J.K., Yariv, L., 2006. Conformity in the lab. Working paper, Caltech.
Grossman, S.J., Kihlstrom, R.E., Mirman, L.J., 1977. A Bayesian approach to the production of information and learning by doing. Rev. Econ. Stud. 44, 533–547.
Halevy, Y., 2007. Ellsberg revisited: An experimental study. Econometrica 75, 503–536.
Hirshleifer, D., Teoh, S.H., 2009. Thought and behavior contagion in capital markets. In: Hens, T., Schenk-Hoppe, K.R. (Eds.), Handbook of Financial Markets: Dynamics and Evolution. North-Holland/Elsevier, pp. 1–46.

Kaivanto, K., Kroll, E.B., 2012. Negative recency, randomization device choice, and reduction of compound lotteries. Econ. Letters 115, 263–267.
Knight, F.H., 1921. Risk, Uncertainty, and Profit. University of Chicago Press, Chicago.
Liu, H.-H., Colman, A.M., 2009. Ambiguity aversion in the long run: Repeated decisions under risk and uncertainty. J. Econ. Psych. 30, 277–284.
Merlo, A., Schotter, A., 1999. A surprise-quiz view on learning in economic experiments. Games Econ. Behav. 28, 25–54.
Mirman, L.J., Samuelson, L., Urbano, A., 1993. Monopoly experimentation. Int. Econ. Rev. 34, 549–563.
Mueller, A., Scarsini, M., 2002. Even risk averters may love risk. Theory Dec. 52, 81–99.
Muthukrishnan, A.V., Wathieu, L., Xu, A.J., 2009. Ambiguity aversion and persistent preference for established brands. Manage. Sci. 55, 1933–1941.
Pulford, B., Colman, A.M., 2008. Size doesn't really matter: Ambiguity aversion in Ellsberg urns with few balls. Exper. Psych. 55, 31–37.
Rasmusen, E., 2010. Career concerns and ambiguity aversion. Econ. Letters 108, 175–177.
Rick, S., Weber, R., 2010. Meaningful learning and transfer of learning in games played repeatedly without feedback. Games Econ. Behav. 68, 716–730.
Rode, C., Cosmides, L., Hell, W., Tooby, J., 1999. When and why do people avoid unknown probabilities? Testing some predictions from optimal foraging theory. Cognition 72, 269–304.
Rosenboim, M., Shavit, T., Cohen, C., 2013. Do bidders require a monetary premium for cognitive effort in an auction? J. Socio-Econ. 42, 99–105.
Savage, L.J., 1954. The Foundations of Statistics. Wiley, New York.
Segal, U., 1990. Two-stage lotteries without the reduction axiom. Econometrica 58, 349–377.
Spears, D., 2012. Poverty and probability: Aspiration and aversion to compound lotteries in El Salvador and India. Exper. Econ., http://dx.doi.org/10.1007/s10683-012-9333-9.
Teodurescu, K., Erev, I., 2011. On the decision to explore new alternatives. Working paper, Technion.
Tversky, A., Kahneman, D., 1992. Advances in prospect theory: Cumulative representation of uncertainty. J. Risk Uncertainty 5, 297–323.