Blindness to the Benefits of Ambiguity:
The Neglect of Learning Opportunities

Stefan T. Trautmann

Tilburg University

Richard J. Zeckhauser

Harvard University

October, 2010

Abstract: Financial, managerial, and medical decisions often involve alternatives whose possible outcomes have ill defined probabilities. In contrast to *risky* alternatives, i.e., those whose probabilities are known, these *ambiguous* 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 marbles drawn from bags, decision makers often prove to be blind to the learning opportunities offered by ambiguous probabilities. Such decision makers violate rational decision making and forgo significant expected payoffs. Decision makers who violate in this fashion fall into four categories: underestimators (of the value of learning), minimum learners (draw insufficient inferences), insensitive stickers (ignore information that they should switch colors), and fallacious switchers (switch colors despite information that they should stick).

Keywords: risk; ambiguity; learning; information; gambler's fallacy; Ellsberg Paradox

JEL codes: D81; D83; G11, C91

In decision making under uncertainty, economists have long distinguished between a situation where the probabilities of outcomes are known (*risk*) and those where probabilities are unknown (*uncertainty*). The distinction was introduced in the works of Knight (1921) and of Keynes (1921). Ellsberg (1961) suggested that people will usually pay a premium to choose an option involving risk rather than one involving uncertainty, even if the two would be considered equivalent under expected utility. Ellsberg employed the more descriptive term *ambiguity* in lieu of uncertainty. We employ his terminology. Numerous studies in the wake of Ellsberg have confirmed a strong aversion to ambiguity, utilizing experiments requiring choices between ambiguous and risky options. The typical study presents the risky options and the ambiguous options in such a manner that rational subjects should be indifferent to the distinction (Raiffa 1961).

Ambiguity avoidance can be justified as a rational response to a range of situations. In the real world, ambiguous options are often offered to people by others whose interests are mostly opposed to theirs. If so, the ambiguous probability will on average disappoint. For example, if a salesman provides information which, if viewed merely as impartial data, would indicate that the breakdown likelihood for a product is about 10%, a strategic assessment would put that breakdown likelihood higher. The salesman is surely selective in what he reveals, what product he chooses to sell, or both (Akerlof 1970; Frisch and Baron 1988; Morris 1997).[1]

Such rationalizations of behavioral biases, whether they come from folk stories or from misplaced extrapolations from slightly dissimilar situations, are reasonable. However, they also suggest that individuals should *recognize the benefits* inherent in commonly occurring ambiguous situations. In contrast to the typical one-shot decision in Ellsberg experiments, ambiguity is often encountered in real decisions that offer repeat choices. In those, ambiguity on probabilities brings benefits in a wide variety of situations. Perhaps the most important benefit is the potential for learning, which implies an option value for the ambiguous choice.[2] Consider hunters in a primitive setting. They can go to their traditional locale, where there is a 20% chance of bagging a catch in a day. Alternatively, they can try a new locale, where their best guess of the chance of catching something is perhaps 18%, but this value is highly uncertain, and

[1] Haisley and Weber (2010) and Ivanov (2010) provide comprehensive discussions of ambiguity attitudes in interpersonal settings.

[2] For theoretical approaches that capture the intuition that ambiguity offers learning opportunities, see Krueger and Svensson, 2009; Mueller and Scarsini, 2002; and Rasmusen, 2010.

could be as high as 40% or as low as 5%. If this were a one-day-only mission, the traditional locale would be preferred. But hunting for one's sustenance is a repeat-choice situation. If the hunters go to the new locale on a dozen days and catch game on 3 or more of them, they will likely have added a superior hunting locale to their permanent assets.[3]

Similar situations arise frequently in medical practice (Frank and Zeckhauser 2007). A patient is prescribed a drug for high cholesterol. It is successful, lowering her total cholesterol from 230 to 190, and her only side effect is a mild case of sweaty palms. The physician is likely to keep the patient on this drug as long as her cholesterol stays low. Yet, there are many medications for treating cholesterol. Another might lower her cholesterol even more effectively or impose no side effects. Trying an alternative would seem to make sense, since the patient is likely to be on a cholesterol medication for the rest of her life. Far too often, as the experiments in our paper will show, individuals are poor at recognizing that, in repeat-play situations, ambiguity in the payoffs to an alternative, holding the expected first-trial payoff fixed, offers an advantage, not a disadvantage.

Ambiguity, Finance and the Financial Crisis

The financial world is dramatically characterized by ambiguity. Investors, analysts and scholars ply their craft by trying to estimate unknown probabilities. Thus, standard discussions in finance revolve around estimating the variance-covariance matrix of returns, a matrix that is known to be always evolving.

The vocabulary of finance, however, needs a richer set of terminology to allow for a world that sometimes goes out of kilter. As Reinhart and Rogoff (2009) show, this is a not an infrequent occurrence. When severe disorder does happen, probabilities that jostle and evolve in normal times instead carom dramatically as they reverberate in response to unanticipated and extraordinary circumstances. Such was the situation with the financial crisis in commencing in late 2007 and cresting in fall 2008, and to a lesser extent in the dramatic price recovery starting in spring 2009. This was not a world characterized by mere ambiguity. Rather, we might think of it as characterized by ambiguity squared.

[3] We are leaving aside the possibility that game moves, responds to hunting, etc., which would probably enhance the benefits of trying a new hunting ground.

In an ambiguity squared world, people are very insecure about what they know. Consequently, small slivers of information can cause prices to leap and plummet, with decisions made on a far from Bayesian basis. And so it was in the fall of 2008. Volatility was enormous. Prices moved dramatically, often in response to little information about the real economy.

Individual decision makers fare poorly, our analysis will show, in such a world. First, in such unfamiliar circumstances, with significant amounts of money being gained, or more often lost every day, anxiety will be pervasive. When anxious, individuals have great difficulty responding in rational fashion to new information. Indeed, when anxiety is high, it is not merely that individuals do not update information with their probabilistic assessments. They tend not to even incorporate probabilities into their thinking.

Recent work demonstrates the major impact of ambiguity in financial decisions. Charness and Gneezy (2010) show experimentally that people are willing to pay a price to avoid ambiguity, and that it affects portfolio composition. Easley and O'Hara (2009) demonstrate that such ambiguity attitudes can affect market prices and provide a rationale for regulation in financial markets.

In this paper, we want to make a more fundamental argument about ambiguity and probability estimation: Few individuals recognize that ambiguity offers the opportunity for learning. If a choice situation is to be repeated, ambiguity brings benefits, since one can change one's choice if one learns the ambiguous choice is superior.

<u>Blindness to the Benefits of Ambiguity</u>

We test whether people understand the benefits of ambiguity, using a variety of repeat-play experiments. In essence, we determine whether individuals choose to play in ambiguous as opposed to risky situations offering the same initial expected probabilities. In these situations, rational Bayesian analysis would clearly show that in each case the ambiguous situation offers substantially higher expected value, and that it stochastically dominates its risky comparator. Yet very few subjects see this advantage; very few go with ambiguity. Consider some consequential real world situations, all related to finance, where the benefits of ambiguity appear to be neglected at present. In the summer of 2010, the consensus estimate is that there are five applicants for every job opening, yet major employers who expect to hire significant numbers of workers once the economy turns up are sitting by the sidelines and having current workers do overtime. The favorability of the hiring situation is unprecedented in recent years. Thus, it would seem to make sense to hire a few workers, see how they perform relative to the norm. If the finding is much better, suggesting that the ability to select in a very tough labor market and among five applicants is a big advantage, then hire many more. This situation, where the payoff to the first-round decision is highly ambiguous, but perhaps well worthwhile once learning is taken into account, is a real world exemplar of the laboratory situations investigated in this paper. In this context, blindness to the benefits of ambiguity produces a much slower recovery to the financial meltdown and subsequent economic slump than might be possible.

It is widely asserted that leading up to the financial crisis even sophisticated investment firms failed to conduct due diligence on assets that subsequently proved to be worth much less than their market price. A common explanation is that they were herding, or thought they were free riding on the information of other players, who were doing much the same thing.

This situation too can be thought of as one where the learning opportunities from ambiguity were missed. It may well have been that detailed investigation of any single ambiguous asset would not be worthwhile on an expected value basis. But let us posit that the firm thought there was a 5% chance that the asset would be found to be shaky. The positive expected payoff would come not from avoiding the purchase of that single asset, but avoiding many dozens of others that were likely to be afflicted with the same shakiness. Indeed, once a class of assets was discovered to be at an unexpected risk level, it would have been possible to go short rather than long. In financial markets, where mars in the quality of one asset tells us about the diminished value of another, securing returns from unrecognized information is a scalable strategy. If every firm did this, of course, the returns would vanish. It appears that virtually no firm engaged in such investigation. Indeed, our results support the prediction that people may rather look for possibly uninformative social evidence rather than to pursue learning opportunities themselves (Goeree and Yariv 2006; Hirshleifer and Teoh 2009). Such learning opportunities are commonly overlooked, making social observation the prominent factor to consider. Had merely a dozen significant firms capitalized on the learning benefits from ambiguity, the bubble would have been much smaller, and losses much more contained.

The current authors hypothesized that the phenomenon of *blindness to the benefits of ambiguity* applies broadly, and that it can be identified in choices made under laboratory conditions. We conducted a set of experiments building on the classic Ellsberg design. The important innovation was that we allowed for repeat choices and, therefore, for the opportunity for beneficial learning. We find that blindness to learning opportunities under ambiguity is deeply rooted. It is found even in these experimental settings that control for outside factors, factors that might intrude on the decision. Such blindness favors the status quo, since a subject would not learn the advantage of switching when it existed.[4]

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 has a different focus than these studies: learning is an inherent part of the decision problem. Consistent with the previous literature, however, we also find that there is little improvement and transfer among tasks. Our paper is closest to those of Charness and Levin (2005) and Charness, Karni , and Levin (2007), which show defects in learning if Bayes' rule and reinforcement learning give different predictions in decision problems whose probabilities are known. We address the basic learning paradigm when some probabilities are not known; that is under ambiguity. We observe that people's lack of a clear understanding of learning under ambiguity leads them to adopt non-Bayesian rules. In this sense our paper shows that the violations identified under risk apply more broadly under ambiguity.

Investigating the nature of the learning defect, we find four different types who violate principles of rational learning from ambiguous situations. We call these four groups underestimators, minimum learners, insensitive stickers, and fallacious switchers. Those in the first two groups underestimate learning opportunities although they get the direction of inference correct. Insensitive stickers and fallacious switchers do worse; they make choices against the inferences from learning.

According to Tolstoi, happy families are all alike, while every unhappy family is unhappy in its own way. A similar observation seems to hold true for situations involving ambiguity: There is only one way to capitalize correctly on learning opportunities under ambiguity, but there are many ways to violate reasonable learning strategies.

This paper is organized as follows. In the next section, we provide an illustration of the blindness to learning opportunities in a medical setting. In Sections 2, 3, and 4, we show failures

[4] The Ellsberg paradox has sometimes been interpreted in terms of difficulties with compound lotteries and violations of the reduction axiom (Halevy 2007, Halevy and Ozdenoren 2008). The rationale put forward in the current paper also applies under the compound lotteries interpretation.

to learn in rational fashion under ambiguity in three incentivized lab experiments, and we identify the four groups of violations discussed above. Section 5 studies the limits of the learning violations, and Section 6 discusses the implications of the results and concludes.

1. Experiment 1: Learning Opportunities in Medical Situations

To provide a first test of our claim that people might not recognize or appreciate the learning opportunities that real-world ambiguous decisions usually provide, we presented subjects with a simple medical decision. Participants were asked to advise a friend who had been presented by his general practitioner with the choice between two drugs to treat a chronic migraine. The first drug had, for each migraine that the person suffered, a 50% chance of mitigating the pain. It was not clear whether this was a good success rate. The second drug, which worked much better for some people and much worse for others, on average also offered a 50% chance over all people. The learning opportunity for the more ambiguous medication lay in the fact that trying the second drug for a few weeks might reveal that the friend belonged to the group of people for whom it worked much better and that he could then successfully use this drug for many years. On the downside, it might not work as well as the other drug. In that case, he might suffer unnecessary pain for some time, but then could switch back to the drug that worked in 50% of the attacks.

To manipulate the potential downside cost, we had the friend obtaining a prescription before a trip that would take either three months or ten days. In the former case, upon recognizing after a month that the new medication did not work well, he could not switch back immediately and would be locked in with the poor medication for some more time. Note that the potential benefit – switching to a superior drug for the rest of one's life (or at least until something better came along) was not affected by this manipulation. In the second case, the ten-day period was arguably too short to affect the downside because a number of migraines were necessary to evaluate the medication. Since the testing period would extend beyond the trip, there was effectively no period of regrets for that case. We also included a condition in which there was no trip and the friend was simply offered the choice as he contacted his doctor for a new prescription. This condition was included to test whether the mere fact that a trip was involved triggered feelings of caution and of potential for blame that might have guided people away from trying the new drug.   Results are given in Table 1.

Table 1. Choice of Migraine Prescription

	3-month trip	10-day trip	No trip
N	63	63	28
Ambiguous chosen	27%	29%	25%
binomial test, two-sided	$p<0.001$	$p<0.001$	$p=0.01$

Notes: 154 undergraduate students were recruited on campus and participated in a computerized questionnaire that included the medical problem. The questionnaire took about 30 minutes, and participants received a €4 flat payment. Each person participated in only one condition.

We took as the null hypothesis that individuals merely chose at random between the two drugs. Rational learning would lead the individual to select the ambiguous alternative for the No trip alternative, and possibly for either of the other two.[5]

Clearly, subjects avoided the ambiguous option more than random or rational choice would prescribe. Moreover, there was no difference among the three trip treatments. This suggests that the 70+% of individuals who stuck with the first drug across treatments were unlikely to engage in a trial later; they seemed to be simply letting ambiguity aversion swamp any learning benefits, quite possibly because they were blind to learning.

The remainder of this paper conducts a series of experiments designed to tease out the various types of violations of rational learning behavior. Each experiment has an ambiguous option that provides for formal learning in a laboratory decision task. Moreover, the benefits of such learning in each are subject to mathematical calculation. Wed did not expect any such calculations by our subjects, but they do permit calculations of loss in expected value.

[5] The 10-day trip would allow for the beginning of learning. On the 3-month trip, the ambiguous alternative would have a downside for a subject who experiences increasing marginal costs to unmitigated migraines.

2. Experiment 2: Repeated Ellsberg Choices

2.1 Design

Ninety-nine undergraduate students participated in a laboratory experiment with real monetary incentives. To model a learning opportunity under ambiguity, we conducted Ellsberg two-color choice experiments. In a baseline condition (BASE), we replicated the classic Ellsberg one-shot decision task, offering subjects a simple choice between a risky bet and an ambiguous bet. The risky bet involved betting on the color of a poker chip drawn from a bag known to contain 5 red and 5 black chips. If the subject correctly guessed the color of the chip, she won €10; otherwise she won nothing. The ambiguous bet involved betting on the color of a chip drawn from a bag with 10 chips, either red or black, in an unknown proportion. 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 risky bag, and then, while wearing a blindfold, put 10 chips into the ambiguous bag. If not stated otherwise, this procedure was followed in all Ellsberg-choice tasks in this study. The transparent procedure was chosen to minimize any suspicion and to emphasize the expected symmetry of colors in the ambiguous option. It follows from the design that, in this baseline one-shot decision, the expected winning probability equaled 50% for both the risky and the ambiguous options, and the expected payoff from either option, therefore, equaled €5.

To introduce the potential for learning in this setting, we included the following repeated Ellsberg-choice task (REPEAT), in which the subjects chose between betting twice on the color of a chip drawn from the risky bag or from the ambiguous bag. The task had two stages. In the first stage, the subject bet on the color drawn from the bag she had chosen, and was paid according to her prediction, as described above. By playing the first stage, she automatically observed the color of a chip drawn from the bag she chose. After the first stage, the chip was replaced in the bag. The subject then played the game again, making a new prediction for another €10 prize, predicting a different color than in the first stage if she chose to. Subjects could not switch between bags after stage 1: they had to choose one of the bags for the complete two-stage game in advance. The repeated game structure was explained to the subjects and illustrated schematically by a time line before they made their choices between the risky and the ambiguous option.

In the repeated game, for the risky option the probability is also 50% in each stage. This is also true for the first stage of the ambiguous option. In the second stage of the ambiguous option, however, it is possible to increase the chances of winning because of the learning opportunity provided by the first stage. For instance, the drawing of a red chip in the first stage would be recognized by a rational subject as an indication that red is more likely. Specifically, by predicting the color drawn from the ambiguous bag in the first stage also for the second stage, the subject achieves a winning probability in stage 2 of 54.5%.

While in the basic Ellsberg experiment both options are equally good from a Bayesian perspective, in the repeated game the expected value of the ambiguous option is €10.45 versus €10 for the risky option. Similarly to Experiment 1, we also wanted to test the effect of a variation of the relative benefit of learning opportunities. We, therefore, included a second repeated Ellsberg choice treatment (REPEAT4), with exactly 2 red and 2 black chips in the risky bag and four chips, either red or black but in an unknown proportion, in the ambiguous bag. Otherwise the choice task was identical to the 10-chip, repeated, Ellsberg-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 risky option is again €10, while for the ambiguous option the expected payoff now equals €11.20, a 12% increase in expected payoff from correctly choosing in Stage 2.[6]

2.2 Results

Table 2 shows the results. In the baseline one-shot treatment, we replicated the typical pattern of ambiguity aversion, with only 23% choosing the ambiguous option. No ambiguity preference was observed in the repeated play treatments with the strictly better ambiguous option due to learning. The ambiguous option received less than 50% of the play in both the 10-chip and the 4-chip treatments, but only in the 4-chip treatment could we reject the hypothesis that people chose randomly. No difference between treatments was significant (χ^2 tests, all $p>0.10$), implying also that the manipulation of the learning benefit between the 4-chip and the 10-chip treatments had no effect. The 4-chip urn elicited even slightly fewer ambiguous choices, providing strong evidence against people's properly recognizing the learning opportunities. To be fair, it may be

[6] For the simple baseline game with no learning opportunities, it has been shown that the size of the Ellsberg urn (that is, the number of chips) does not affect behavior (Pulford & Colman 2008).

harder to recognize that the 4-chip situation offers more learning than the 10-chip situation, than to recognize that drawing a chip does enable one to learn.

Table 2. Choice of Ambiguous Ellsberg Bet

	BASE	REPEAT	REPEAT4
N	31	33	35
Ambiguous chosen	7 (23%)	13 (39%)	10 (29%)
binomial test, two-sided	p=0.003	p=0.296	p=0.017

Table 3 provides more evidence regarding the neglect of learning benefits under ambiguity. In this table, we show for the repeated-bet treatments the number of subjects who stayed with their first-stage prediction after either a successful or an unsuccessful prediction. Of the subjects who chose ambiguous, only about 60% behaved consistently with learning. Thus, there were subjects who stuck with their initial prediction, despite choosing a color contrary to it, although this sampling outcome signaled that the other color was more likely. This group included about 22% of the ambiguous choosers. Their behavior may either derive from a strong gambler's fallacy belief ("now red is due"), or may be caused by an insensitivity to learning and by a decision to persist with a preferred color. We also observed that 17% of subjects who switched after an initial success, although such a success indicated the initial color was more likely. We label these individuals fallacious switchers. They too fell prey to the gambler's fallacy. These two groups, the the insensitive stickers and fallacious switchers, thus, did not simply underestimate learning opportunities. They either drew wrong inferences, or drew no inferences from new information.

Table 3 also shows that subjects who chose the risky option sometimes switched after an initial success, but almost never switched after an initial failure. That is, they deviated from purely random choice following some strategy. It seems that, under ambiguity, such strategies played a role as well, although they violate optimal learning behavior. This can also be seen from the comparison of learning-compatible choices under ambiguity with "as-if" learning-compatible choices under risk. "As-if" learning indicates a choice that would have been compatible with learning had the subject been playing the ambiguous option. While there is no learning for the risky option, this measure allows us to identify choice patterns for comparison with the behavior under ambiguity. Pooling data from the 4-chip and the 10-chip games, we can only marginally reject the hypothesis that ambiguous choosers make more learning-compatible choices than the risky choosers (χ^2 test, all p=0.073, one-sided). Clearly, they do not make better than random choices, either.

Table 3. Learning Behavior in Repeated-Bet Game

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

Note: * "As-if" learning for risky option: indicates a choice that would have been compatible with learning had the subject played the ambiguous option

To our knowledge, there exists no experimental evidence today on repeated Ellsberg experiments that offer learning opportunities. Yet, we believe that situation is the norm in the world, albeit not with such artificial situations as drawing a ball from an urn. But if one has a repeat choice of dry cleaners, driving routes, restaurants, or employees, where their variability in performance is to be expected, there is 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 locales. Few individuals recognize the two-armed bandit feature of the problem, and the advantage of gathering information on Route B.

We should note that some past experiments involving no learning possibility have shown a tilt toward ambiguity. Static repetition from newly assembled urns and repeated bets without switching opportunities has been studied by Liu and Colman (2009). If the prize is larger for the ambiguous urn and 100 repetitions are made from the ambiguous urn, subjects prefer the ambiguous. These authors argue that, through repetition, the uncertainty in ambiguous situations is reduced to a 50% chance as in the risky situations, 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. They found that ambiguity is preferable from a profit maximization point of view. Indeed, they found that people prefer ambiguity if x is large, implying that the risky situation, therefore, gives a low chance of achieving x successes.

Our results replicate basic patterns observed by Charness and Levin (2005) in decisions that involve only known probabilities. Optimal switching strategies are violated by a large group of subjects, 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 produce a much stronger bias against the favorable urn. 60% to 70% of the subject chose the unfavorable urn. In their study, a pure risk setting, only 20% to 30% of subjects chose the unfavorable urn.

3. Experiment 3: Making Learning Opportunities Salient

In Experiment 2, we identified learning errors, in particular the strong gambler's fallacy and the insensitivity to learning-relevant information of the subjects choosing the ambiguous option. (Subjects choosing the risky option had no opportunity of falling into such error, since they learned nothing.) Furthermore, the learning behavior emerged endogenously from observations of outcomes in the ambiguous option, similar to how learning in ambiguous situations emerges in real world decisions. It might have been that participants overlooked the learning opportunity. While this is also likely the case in many real-world problems, we conjectured that even if the possibility of learning were made salient, people might not perceive strong benefits from learning in situations in which uncertainty could not be resolved. To test this, we conducted an experiment in which people were forced to learn before making their choices between risky and ambiguous options, and where we could identify learning also for those who chose the risky option.

3.1 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, with the second part being identical to the 10-chip Ellsberg two-color bet for a prize of €10 in BASE. Before making their choices, in the first part of the experiment subjects had to draw one chip from the risky bag and one chip from the ambiguous bag, always with replacement. They noted the colors sampled, and then had to predict the contents of the urn, that is, whether they expected more red, more black, or an equal number of red and black chips in each urn. If both predictions 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 Ellsberg choice in Part II. At the end of the experiment, one of the two parts was chosen by coin toss for real payment.

For the ambiguous bag, sampling a red chip implied that the bag was more likely predominantly red than predominantly black or equally distributed. For the risky 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 risky bag.

Sampling from the ambiguous bag in Part I of the experiment forced the subjects to learn about the distribution of colors, and to upgrade their expected likelihood of winning the prize for the Part II Ellsberg choice. 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 repeated Ellsberg choice. In the repeated game, the subjects would still have to make the first bet without any information; hence, the effect on total expected earnings would be smaller. Because all the subjects had to sample the ambiguous bag and make a prediction regarding its contents, we could also observe learning errors for those choosing the risky option.

3.2 Results

Table 4 shows the results. We found a similar level of ambiguous choices as in the previous experiments, with 36% choosing ambiguous (p=0.079, binomial test, two-sided). Overall, in the prediction of the contents of the ambiguous urn, 26% violated learning and did not predict a majority of the color they drew. The incidence of this violation was similar among risky and ambiguous choosers. For ambiguous choosers, we also observed whether they violated learning 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.

Table 4. Predicted Composition and Ellsberg Choices after Forced Learning

	Part II Ellsberg Choice		
	Risky	Ambiguous	All
Ellsberg Choice	30 (64%)	17 (36%)	47
Part I: violate learning in prediction of bag composition	7 (23%)	5 (29%)*	12 (26%)
Part II: violate learning in color prediction	–	4 (24%)*	
Total incidence of learning mistakes	7 (23%)	7 (41%)	14 (30%)

Note: * These two groups had two members in common.

Among the 35 subjects who stated that they believed that the color they picked in Part I in the ambiguous bag was in the majority, 23 (roughly two thirds) then chose the risky bet in Part II. We asked these subjects why they had not chosen the ambiguous option, and 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 that they had predicted the majority composition more or less randomly. Since they had seen no clear evidence for either color, they had preferred the risky option in the Part II bet. However, as the data show, it was still more natural for the subjects to predict the contents according to their samples. Vice versa, drawing a red chip seemed not to be perceived as evidence for a majority of black. While this group might basically have had the right intuition, they underestimated the value of the sample. Another group of subjects, who fall into our class of minimum learners, announced the learning-compatible contents but then chose the risky option, suggested a maximin way of thinking: 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 risky bet, because there were assuredly 5 red chips in the risky option versus assurance of only 1 red chip in the ambiguous option.[7]

4. Experiment 4: Putting a Price Tag on Benefits from Learning

The previous experiment replicated the strong violations of learning that we found in Experiment 2, but it also suggested that those who do not commit such strong violations may not perceive the clear benefits of ambiguity either. To study how well people are calibrated when learning under ambiguity, we designed an experiment in which the subjects would not choose between risky and ambiguous urns. Rather, they would only make bets on ambiguous options, some with learning and some without learning.

4.1 Design

Forty-three subjects participated in an experiment in which they had to predict the color of a chip drawn from an urn with 4 chips, either red or black in an unknown proportion.[8] Without any sample, the probability of winning the price in this bet equals 50%. With a sample of one chip with replacement, predicting the color sampled increases the chance of winning to 62% and the expected value of the gamble by 24%. Note that the percentage benefit from learning is larger in this experiment than in the repeated Ellsberg choice with 4 chips. 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 ambiguous 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

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

[8] This experiment used marbles instead of poker chips, and these marbles where drawn from a box of 25 marbles of each color, instead of 50 of each, as in the previous experiments. Calculations account for the size of the box. For consistency in presentation, we refer to “chips“ in the text.

winning prize of €x. There existed some x<20 such that the subject was 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 in a bag 39 slips of paper with prize offers between €.50 and €19.50 in equal steps of €.50. Subjects wrote down their LAP. If the randomly selected offered prize "y" was 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 €16.2 because winning €16.20 with a 62% chance offers an expected value equal to winning €20 with a chance of 50%, in the case of no sampling opportunity. The optimal strategy, obviously, involves betting on the color sampled.

4.2 Results

Based on the optimal LAP of €16.20, we defined people as well-calibrated learners 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 it includes the prominent amount of €15.[9] Thus, we applied a conservative criterion for learning neglect here. Table 5 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). Indeed, the average LAP in the last group was below the expected value of the ambiguous bet in the absence of sampling opportunities. This group seemed to hold the maximin view discussed above: that felt they learned very little, but compared to the situation of complete ignorance, they strictly preferred one sample.

[9] Most people specified full Euro amounts.

Table 5. Valuation of Learning Opportunity

		Well calibrated 14≤LAP≤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 distribution within row

The remaining rows of Table 5 distinguish subjects according to their betting strategies. Individuals who employed the correct learning strategy had a mean LAP about €15. They were well calibrated 58% of the time.

There were two groups who followed strategies contrary to the inference from the first draw: they bet against the color drawn. Thus, 19% of the sample got the first draw wrong, but stuck with their color. We call them insensitive stickers. 16% of subjects got the first draw correct, but then switched colors. We label them fallacious switchers. Interestingly, both of these groups had lower average LAPs, respectively €13.50 and €11.29 than did those following correct strategy. Thus, they thought they were learning a lot. Both of these groups were following the gambler’s fallacy: assuming that a color had not come up the first time was more likely to on the second trial. Such choices are well documented for risky choices, 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. However, some optimal bettors were surely making choices independent of their sample draw. Thus, they might have chosen their bet at random, or just stuck with or just switched colors independent of the first marble drawn. Similarly, the 15 subjects whose bets violated learning were assuredly choosing absent learning. Assuming the nonlearners divided evenly, there would be 15 lucky bettors on the right color, implying 13 correct bettors – less than one third – who were relying on learning.

5. Experiment 5: The Limits of Learning Neglect

We have shown in the previous experiments that various failures to learn under ambiguity exist in situations where 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 and ambiguity would be taken by the subjects. Second, we predicted that experience in a learning task that revealed the general principle of learning under ambiguity would affect decisions in a task where learning was less obvious to people, as we have shown above.

5.1 Design

Thirty-two subjects participated in an experiment that consisted of two parts, with each part presented as a separate experiment to the subjects. Each part involved monetary incentives; at the end of the experiment one part was randomly selected by coin flip for real payment. The second part of the experiment was identical to the 4-chip repeated Ellsberg choice experiment presented in Section 2. The first part of the experiment involved the following choice situation, modelled 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 and which contained only black marbles. Before learning about the decision problem, subjects chose one of the blue bags (and the white bag) and the other blue bag was removed.

The decision problem was similar to the learning setting in Section 2. The subjects had to choose either the white bag with the known and equal distribution of red and black, or an ambiguous blue bag containing only either two red marbles or two black marbles, to make a repeated bet on the color of a marble drawn, with replacement, from the bag. 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 risky white bag in both the first and the second drawings. For an ambiguous blue bag, the first drawing also offered a 50% chance of winning; 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 ambiguity of Part I, the blue bag offered a certain gain of €5 in the Stage II. After the first stage was finished, the bags were set aside; then Part II, the repeated Ellsberg, was conducted exactly as described in Section 2.

5.2 Results

Table 6 shows the results. In the Part 1 repeated-betting task with complete resolution of uncertainty for the ambiguous option in the second drawing, 75% of the subjects chose the ambiguous option ($p=0.007$, binomial test, two-sided). That is, in contrast to the previous experiments, the majority correctly understood the learning opportunity available under ambiguity. Given the simplicity of the task, it is somewhat surprising that there were still subjects who did not see the benefit of ambiguity. However, we have observed serious learning failures before. 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 a demonstration of the complete resolution of uncertainty can reduce but hardly eliminate learning neglect. For subjects who intuited the basic advantage but underestimated the learning potential, and for those who applied maximin learning, this task's structure immediately revealed the benefit of ambiguity.

Table 6. The Limits of Learning Neglect

Part 1:	Risky (white: 1 red and 1 black)		Ambiguous (blue: either two red or two black)		All	
	8 (25%)		24 (75%)		32	
Part 2:	Risky	Ambiguous	Risky	Ambiguous	Risky	Ambiguous
	7 (88%)*	1 (12%)*	15 (63%)*	9 (27%)*	22 (69%)*	10 (31%)*
“As-if” learning	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 Ellsberg 4-chip task. “As-if” learning for risky option: indicates a choice that would have been compatible with learning had the subject played the ambiguous option. * Probabilities conditional on Part 1 choice.   ** Probabilities conditional on Part 1 and Part 2 choices.

While the subjects successfully identified the learning opportunity in Part 1, only 31% chose the ambiguous option in the repeated Ellsberg 4-chip task (p=0.050, binomial test, two-sided), reversing the preference demonstrated in Part 1. The percentage of people choosing the ambiguous option was very similar to that in Experiment 2. Clearly, the Part 1 experience did not successfully train subjects to understand the basic concept of the learning benefits under ambiguity. While most subjects who chose the risky option in Part 1 also chose the risky option in Part 2, 63% of those who chose the ambiguous option in Part 1 switched to the risky option in Part 2. There were observed effects of the Part 1 experience in terms of optimal learning. The bottom row of Table 6 shows the number of subjects who behaved in accordance with learning (or “as-if” learning for risky choosers, as defined in Section 2) in the repeated Ellsberg task. In contrast to our results in Experiment 2, in which many ambiguity-choosers violated learning, here all the ambiguous choosers correctly chose for the second drawing according to the color observed in the first draw. Interestingly, the data on “as-if” learning by risky choosers revealed that these people had also likely been influenced by the preceding statistical decision task in Part 1. While actual betting behavior did not affect the chances of winning for the risky option, the subjects must have tried to use a strategy. Specifically, the only way to reduce “as-if” learning in the risky option to significantly below 50% (which is the case here: p=0.017 binomial test, two-sided), was adverse betting by switching to the color not extracted in the first draw. While not harmful in terms of payoffs in that task, the experience may have led to considerably fewer preferred strategies and outcomes in other tasks. Clearly, after experiencing the Part 1 task before making the repeated Ellsberg 4-chip choices, the ambiguous choosers did better than random (p=0.050, binomial test, two-sided) and better than the risky choosers in terms of learning or “as-if” learning (χ^2 test, all p=0.022).[10]

6. Discussion and Conclusion

This paper contributes two important insights regarding real world decisions under ambiguity. First, ambiguity about success probabilities will in many cases be accompanied by a learning opportunity and by future opportunities either to make the same choice again, or to switch choices. Second, real world decision makers seem oblivious to such learning opportunities; thus ambiguity aversion dictates their decisions]. Researchers of decision making have also sometimes rejected the normative benefits of learning. Schneeweiss (1999) refers to the potential short-term losses and overlooks the potentially large long-term benefits from learning. Frisch and Baron (1988) highlight the potential benefit of waiting for new information as a reason to avoid ambiguous alternatives. Such analyses overlook the obvious point that an important way to obtain this new information, (regarding a scalable investment, a new variety of crop, or a new medical treatment, etc.) is often to invest in such an ambiguous alternative.

In our paper we have shown that there is only one way to correctly understand the potential for learning under ambiguity, but many ways to fail to understand it. We identified four groups who violate rational learning: underestimators, minimum learners, insensitive stickers and fallacious switchers. Our data suggest that most individuals fall into at least one of these categories. Thus only a minority of individuals has the potential to make a rational tradeoff between the benefits of ambiguity and other factors, like anticipated blame from bad outcomes (Frisch and Baron 1988; Muthukrishnan et al. 2009).

From an evolutionary perspective, why would learning avoidance persist if the benefits from learning are large? 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 have virtually no pedagogic power (Baumeister et al. 2001). In the current setting, ambiguous options would need to be sampled

[10] These tests obviously have low power because of the few people choosing the ambiguous option.

repeatedly in order to obtain sufficient information on whether to switch from the status quo. Both bad and good outcomes would be experienced along the way, but only good ones could trigger switching. Bad outcomes would also weigh much more heavily, leading people to require too much positive evidence before shifting to ambiguous options. 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 addition, if one does not know what returns would have come from an ambiguous alternative, one cannot feel remorse 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 experience an even higher asymmetry from bad and good payoffs (Eriksen and Kvaloy 2009, 2010). Most people, for that reason, have had many fewer positive learning experiences with ambiguity than rational sampling would provide.

Our results may also add to the understanding of herding behavior and behavioral contagion in financial markets (Hirshleifer and Teoh 2009). In a recent experimental study, Goeree and Yariv (2006) let subjects choose between an informative signal in an ambiguous situation similar to the ones studied in our paper, and an uninformative social signal. Specifically, subjects had to predict the contents of a jar that was filled 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%, and subjects could sample once with replacement from the jar before making their guess (informative statistical signal). Alternatively, they could choose to receive an uninformative social signal, after observing the predictions of 3 people who had randomly guessed the distribution without any statistical information. 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 exists. Our results suggest that even a weak preference for conformity may be enough in the Goeree and Yariv learning paradigm, reconciling their results with those of

Corazinni and Greiner. We showed that most people have no proper concept of learning in situations where uncertainty cannot be 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.

Whether in financial, medical or other decisions, learning opportunities in which outcomes and probabilities are ambiguous offer large gains over known risks. In the spirit of Tolstoi's claim about happy and unhappy families, we have shown that there is only one way to capitalize correctly on these benefits, while there are many ways to violate reasonable learning strategies. From a normative perspective, speaking with Dante now, we will have to go through various circles of erroneous thinking before succeeding at making people recognize learning opportunities. A successful experience in a similar setting proved not powerful enough to make people realize the benefits of ambiguity. Such an experience, however, did lead to fewer mistakes by individuals forced to try out an ambiguous option, suggesting that there may be ways to reduce learning neglect in decision makers.[11]

Ambiguity always lurks in the financial world. At times, as it did in the financial crisis, it emerges as a dominant force. Ambiguity clouds decision making, making crises and bubbles not surprising phenomena.

[11] Though we used strong financial incentives and manipulated experience, we did not study whether team decision making or consulting with others improves performance. Charness et al. (2007, 2010) and Kocher et al. (2006) show such improvements in statistical reasoning tasks. Keck et al. (2010) show a tendency for groups to become more ambiguity neutral, suggesting higher rationality and possibly better learning.

References

Abdellaoui, M., Bleichrodt, H., Paraschiv, C., 2007. Loss aversion under prospect theory: A parameter-free measurement. *Management Science* 53, 1659–1674.

Akerlof, G. A., 1970. The market for “lemons”: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics* 84, 488–500.

Baumeister, R. F., Bratslavsky, E., Finkenauer, C., Vohs, K. D., 2001. Bad is stronger than good. *Review of General Psychology* 5, 323–370.

Becker, G. M., de Groot, M. H., Marschak, J., 1963. Stochastic models of choice behavior. *Behavioral Science* 8, 41–55.

Charness, G., Gneezy, U., 2010. Portfolio choice and risk attitudes: An experiment. *Economic Inquiry* 48, 133–146.

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. *Journal of Risk and Uncertainty* 35, 129–148.

Charness, G., Karni, E., Levin, D., 2010. On the Conjunction Fallacy in Probability Judgment: New Experimental Evidence regarding Linda. *Games and Economic Behavior* 68, 551–556.

Charness, G., Levin, D., 2005. When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect. *American Economic Review* 95, 1300–1309.

Corazinni, L., Greiner, B., 2007. Herding, social preferences, and (non-)conformity. *Economics Letters* 97, 74–80.

Easley, D., O’Hara, M., 2009. Ambiguity and nonparticipation: The role of regulation. *Review of Financial Studies* 22, 1817–1843.

Ellsberg, D., 1961. Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics* 7, 643–669.

Eriksen, K. W., Kvaloy, O., 2009. Do financial advisors exhibit myopic loss aversion? *Financial Markets and Portfolio Management*, forthcoming.

Eriksen, K. W., Kvaloy, O., 2010. Myopic investment management. *Review of Finance*, forthcoming.

Frank, R. G., Zeckhauser, R. J., 2007. Custom-made versus ready-to-wear treatments: Behavioral propensities in physicians' choices. *Journal of Health Economics* 26, 1101-1127.

Frisch, D., Baron, J., 1988. Ambiguity and rationality. *Journal of Behavioral Decision Making* 1, 149–157.

Goeree, J. K., Yariv, L., 2006. Conformity in the lab. Working paper, Caltech.

Haisley, E., Weber, R., 2010. Self-serving interpretations of ambiguity in other-regarding behavior. *Games and Economic Behavior* 68, 614–625.

Halevy, Y., 2007. Ellsberg Revisited: an Experimental Study. *Econometrica* 75, 503-536.

Halevy, Y., Ozdenoren, E., 2008. Uncertainty and Compound Lotteries: Calibration. Working paper, University of British Columbia.

Hirshleifer, D., Teoh, S. H., 2009. Thought and behavior contagion in capital markets. In: *Handbook of Financial Markets: Dynamics and Evolution*. Hens, T., Schenk-Hoppe, K. R., editors. North-Holland/Elsevier, 1–46.

Ivanov, A., 2010. Attitudes to Ambiguity in One-Shot Normal-Form Games: An Experimental Study. Games and Economic Behavior, forthcoming.

Keck, S., Diecidue, E., Budesco, D., 2010. Group decision making under vagueness. Working paper, Insead.

Keynes, J. M., 1921. *A Treatise on Probability*. MacMillan, London.

Knight, F. H., 1921. *Risk, Uncertainty, and Profit*. University of Chicago Press, Chicago.

Kocher, M., Strauss, S., Sutter, M., 2006. Individual or Team Decision making: Causes and Consequences of Self-Selection. Games and Economic Behavior 56, 259–270.

Krueger, N. A., Svensson, M., 2009. The impact of real options on willingness to pay for mortality risk reductions. *Journal of Health Economics* 28, 563-569.

Liu, H.-H., Colman, A. M., 2009. Ambiguity aversion in the long run: Repeated decisions under risk and uncertainty. *Journal of Economic Psychology* 30, 277–284.

Merlo, A., Schotter, A., 1999. A Surprise-Quiz View on Learning in Economic Experiments. Games and Economic Behavior 28, 25–54.

Morris, S., 1997. Risk, uncertainty, and hidden information. *Theory and Decision* 42, 235–270.

Mueller, A., Scarsini, M., 2002. Even risk averters may love risk. *Theory and Decision* 52, 81–99.

Muthukrishnan, A. V., Wathieu, L., Xu, A. J., 2009. Ambiguity Aversion and Persistent Preference for Established Brands. *Management Science* 55, 1933–1941.

Pulford, B., Colman, A. M., 2008. Size doesn't really matter: Ambiguity aversion in Ellsberg urns with few balls. *Experimental Psychology* 55, 31–37.

Raiffa, H., 1961. Risk, ambiguity, and the savage axioms: Comment. *Quarterly Journal of Economics* 75, 690–694.

Rasmusen, E., 2010. Career concerns and ambiguity aversion. *Economics Letters* 108, 175–177.

Reinhard, C. M., Rogoff, K. S., 2009. *This Time Is Different: Eight Centuries of Financial Folly*, Princeton University Press, Princeton.

Rick, S., Weber, R., 2010. Meaningful Learning and Transfer of Learning in Games Played Repeatedly Without Feedback. Games and Economic Behavior 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.

Schneeweiss, H., 1999. Resolving the Ellsberg paradox by assuming that people evaluate repetitive sampling. In: *Modelling and Decisions in Economics: Essays in Honor of Franz Ferschl*. Leopold-Wildburger, U., Feichtinger, G., Kistner, K. P., editors. Springer, 83–95.

Tversky, A., Kahneman, D., 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5, 297-323.