The impact of statistical learning on violations of the sure-thing principle*

Nicky Nicholls† Aylit Tina Romm‡ Alexander Zimper§

February 23, 2015

Abstract

This paper experimentally tests whether violations of Savage's (1954) sure-thing principle (STP) decrease through statistical learning. Our subjects repeatedly had to bet on the drawings from an urn with an unknown proportion of differently colored balls. The control group was thereby subjected to learning through mere thought only. In addition, the test group received more and more statistical information over the course of the experiment by observing the color of the ball actually drawn after each bet. We expected that statistical learning would decrease the decision makers' ambiguity, thereby implying a stronger decrease of STP violations in the test than in the control group. However, our data surprisingly shows that learning by mere thought rather than statistical learning leads to a decrease in STP violations.

Keywords: Learning; Statistical Learning; Sure Thing Principle; Prospect Theory; Independence Axiom
JEL Classifications: C91, D81

---

*We thank Justine Burns, Ulrich Schmidt, Alex Ludwig and ESSA 2013 conference participants for helpful comments and suggestions. We are greatly indebted to Peter Wakker and an anonymous referee for sharpening our theoretical understanding of the experimental findings. Aylit and Nicky gratefully acknowledge financial funding of the experiment by the School of Economics and Business Sciences at Wits. We all gratefully acknowledge financial support from Economic Research Southern Africa (ERSA).

†Department of Economics, University of Pretoria, South Africa. E-mail: nicky.nicholls@gmail.com

‡School of Economic and Business Sciences, University of the Witwatersrand, Johannesburg, South Africa. E-mail: aylit.romm@wits.ac.za

§Corresponding author. Department of Economics, University of Pretoria, Private Bag X20, Hatfield 0028, South Africa. Tel: 0027 12 420-3522. Fax: 0027 86 691-2749. E-mail: alexander.zimper@up.ac.za

Von Neumann and Morgenstern's (1947) (vNM) independence axiom (IA) is central to their axiomatization of expected utility (EU) theory within a framework of decision making under risk. In this vNM framework decision makers have preferences over known probability distributions (i.e. lotteries). Although there exist numerous empirical studies[1] which demonstrate that real life decision makers systematically violate the IA, there also exists empirical evidence which suggests that such violations decline through learning. More specifically, van de Kuilen and Wakker (2006) report a significant decrease in violations of the IA if the respondents experience the consequences of their decisions in terms of played out lotteries, whereas there is virtually no impact on the number of IA violations if the respondents only think repeatedly about their choices.

This paper reports the results of an experimental study which investigates the impact of learning on violations of the sure-thing principle (STP). The STP is central to Savage's (1954) axiomatization of subjective EU theory within a framework of decision making under uncertainty. In this framework, the decision makers do not know any objective probabilities. Starting with Ellsberg (1961), several studies have reported systematic violations of the STP in experimental situations. As a reaction to these empirical studies, subsequent decision theories explain violations of the STP through ambiguity[2] attitudes. For example, prospect theory (Tversky and Kahneman 1992; Wakker and Tversky 1993) and Choquet EU theory (Schmeidler 1989; Gilboa 1987) express ambiguity attitudes through non-additive probability measures whereas multiple priors EU theory (Gilboa and Schmeidler 1989; Jaffray 1994; Ghirardato, Maccheroni, and Marinacci 2004) uses sets of subjective additive probability measures.

Our framework of unknown probabilities allows us to distinguish between two different notions of learning. First, there is, as in the known-probabilities framework of van de Kuilen and Wakker (2006), the notion of *learning through mere thought* according to which the decision makers are repeatedly facing similar choice situations without actually experiencing the consequences of their decisions. Second, there is the notion of *statistical learning* according to which the decision makers may learn probabilities through the observation of statistical information in the form of data generated by these probabilities. Whenever our decision makers receive such statistical information, they simultaneously gain experience by observing the consequences of their decisions. Consequently, our notion of statistical learning includes but does not reduce to van de Kuilen and Wakker's (2006) notion of *learning through experience.*

[1] See, for example, Allais (1979), Wu and Gonzalez (1996), Starmer (2000), Schmidt (2004), Sugden (2004), and references therein.

[2] Somewhat loosely, we speak of "ambiguity" whenever a decision maker cannot comprehensively resolve all his uncertainty through a unique additive–either subjective or objective–probability measure.

Our first conjecture states that violations of the STP should not increase,but should even slightly decrease, through learning by thought. To see why, let us adopt the normative interpretation according to which Savage's (1954) subjective EU theory stands for the rational benchmark model of decision making under uncertainty. Under this interpretation, it seems plausible that intelligent decision makers should not further diverge from rational decision making if they are repeatedly given opportunities to think about their decisions. If anything, after repeated thinking their decision making should become more and not less rational.

In analogy to the findings of van de Kuilen and Wakker (2006), we would expect that statistical learning should strictly decrease violations of the STP because it implies learning by experience. Moreover, observe that the multiple priors approach directly interprets ambiguity as a lack of statistical information. In the words of Gilboa and Schmeidler (1989, p. 142): "[...] the subject has too little information to form a prior. Hence (s)he considers a set of priors." Alternatively, Mark Machina (1994) and Peter Wakker (2010, p. 44) interpret decision making with known probabilities as the limiting case of decision making under uncertainty. In our opinion, it is plausible that decision situations under uncertainty should converge towards this limit through statistical learning. A related formal model of Bayesian updating in which ambiguity vanishes through statistical learning has been put forward by Marinacci (2002).[3] By these interpretations, one would intuitively expect that a decision maker's ambiguity attitudes, and thereby her violations of the STP, will become increasingly irrelevant if she observes more and more statistical information. By collecting the above arguments, our second conjecture states that statistical learning should lead to a further decrease in violations of the STP in addition to any decrease that might have been caused by learning through mere thought.

Our experimental study has put both conjectures to the empirical test. In our experiment respondents were asked to choose between pairs of uncertain prospects such that the actual payout is determined by a ball's color drawn (with replacement) from an urn containing 20 red, 20 yellow and 60 blue balls whereby these proportions were not known by the respondents. In total there were 30 questions, organized as 15 consecutive pairs of questions, such that the answers to each consecutive pair of questions either violate the STP or not. We split the respondents into a test and a control group. After answering

[3] Note that there exist alternative theoretical models of Bayesian learning under ambiguity in which statistical learning does not necessarily reduce the decision maker's ambiguity (cf., e.g. Epstein and Schneider 2007; Zimper and Ludwig 2009; Zimper 2011, 2013). Roughly speaking, in these models ambiguity is, in contrast to Marinacci's (2002) formal argument, not restricted to the parameter-values in question but it also concerns the joint objective probabilities of these parameter-values and the data sample, i.e. the updating process.

each question the respondents of the test group received statistical information in the form of a ball's color drawn from the urn. In contrast, the respondents of the control group did not receive any statistical information as feedback. The control group was thus asked the same questions in the same experimental situation with only the specific experimental treatment—an increase in statistical information—missing. That is, whereas both groups were equally subjected to learning by thought, only the test group was additionally subjected to statistical learning. Our research questions regarding whether STP violations decrease through learning by thought, resp. statistical learning, were thus addressed by the experiment as follows:

1. Does the number of STP violations of the control group decrease over the course of the experiment?

2. Compared to the control group, does the number of STP violations of the test group decrease over the course of the experiment?

At first, we had to assess whether statistical learning, in the sense of subjective beliefs converging to objective probabilities, did actually happen within the test group or not. To this purpose we asked all respondents at the end of the experiment for their estimates about the differently colored balls' proportions. Whereas the control group—having received no statistical information—estimated a roughly equal proportion for all colors, the test group's estimate was very close to the true proportions. We interpret these answers as confirmation that converging statistical learning did indeed happen within the test group while the control group remained ignorant about the true proportions.

Next, turning to our experimental results we continue to see the same prevalence of STP violations in the test group even after multiple rounds of learning. In contrast, a moderate decrease in STP violations is significant for the control group which was subjected to learning through mere thought but not to statistical learning.

On the one hand, our experimental results thus confirm our first conjecture that learning through mere thought implies a slight decrease in STP violations. On the other hand, however, the data does not lend any support to our second conjecture according to which statistical learning should imply an additional decrease in STP violations. To the contrary, our regression analysis points to the possibility that statistical learning might even be responsible for a slight increase in STP violations. This finding is very puzzling to us because it suggests that learning through experience may not have the same impact in ambiguous decision situations as in decision situations with known probabilities. Unfortunately, we do not have any answers yet as to why this might be the case.

The remainder of this paper is structured as follows. Section 1 discusses the experimental methodology used and Section 2 reports the results of the experiment. In Section 3 we discuss the robustness of our findings with respect to our experimental design. Section 4 concludes. An Appendix contains formal arguments as well as a caveat that applies to the interpretation of our experimental design. We provide further details (instructions, questionnaire) of the experiment in a Supplementary Appendix.

1 Experiment

This section describes in detail the experiment that we conducted to test whether violations of the STP decline through learning by thought, resp. statistical learning, or not. As our point of departure consider the following thought experiment of Marinacci (2002):

> "Consider a decision maker (DM) who has to make a decision based on the drawings of an urn of known composition. The confidence he has in his decisions will depend on the quality of the information on the balls' proportion, the more he knows, the more he will feel confident.
>
> Suppose the DM can sample with replacement from this urn before making a decision. Regardless of how poor is his *a priori* information about the balls' proportions, it is natural to expect that eventually, as the number of observations increases, he will become closer and closer to learn the true balls' proportion and become more and more confident in his decisions." (p. 143)

In accordance with Marinacci's thought experiment, we generated uncertainty by means of drawings (with replacement) from an ambiguous urn containing balls of red, yellow, and blue color. Whereas the test group was given the opportunity to gradually learn the true proportions of different colored balls in the urn, the control group did not receive any statistical information over the course of the experiment. The test for violations of the STP follows similar approaches in the literature (e.g. Wu and Gonzalez 1999; Chapter 10.4.3 in Wakker 2010) whereby a caveat applies (see Appendix A4).

1.1 Eliciting violations of the STP

Consider the three events R(ed ball is drawn), Y(ellow ball is drawn), and B(lue ball is drawn). We test for violations of the STP through a choice between prospects $\mathbf{A}$ and $\mathbf{B}$, on the one hand, and a choice between prospects $\mathbf{A}'$ and $\mathbf{B}'$ on the other hand. These pairs of prospects have the following payoff structure:

	R	Y	B
A	y	x	z
B	y	y	y

	R	Y	B
$\mathbf{A}'$	x	x	z
$\mathbf{B}'$	x	y	y

where $z > y > x$ are monetary amounts. In the experiment any such pair of prospects will correspond to a pair of observed choices whereby we interpret these choices as revealed (strict) preferences; that is, we interpret, e.g. the choice pair $\mathbf{A}, \mathbf{A}'$ as revealed preferences $\mathbf{A} \succ \mathbf{B}$ and $\mathbf{A}' \succ \mathbf{B}'$.

In Appendix A1, we formally show that the choice pairs

$$\mathbf{A}, \mathbf{A}' \text{ and } \mathbf{B}, \mathbf{B}' \tag{1}$$

are consistent with the STP, whereas the choice pairs

$$\mathbf{A}, \mathbf{B}' \text{ and } \mathbf{B}, \mathbf{A}' \tag{2}$$

violate the STP. Further, suppose that STP violating decision makers can be described as *prospect theory* decision makers (cf. Wakker 2010; Abdellaoui et al. 2011). Prospect theory in our sense encompasses Gilboa's (1987) Choquet EU axiomatization as well as the axiomatizations of cumulative prospect theory in Tversky and Kahneman (1992) and in Wakker and Tversky (1993) and it has the descriptive advantage that it can simultaneously accommodate violations of the STP as well as of the IA.[4] We demonstrate in Appendix A2 that the choice pair $\mathbf{A}, \mathbf{B}'$ is associated with *optimistic* whereas $\mathbf{B}, \mathbf{A}'$ is associated with *pessimistic* ambiguity attitudes of prospect theory decision makers.

1.2 Subjects

Undergraduate commerce students at the University of the Witwatersrand (Wits) were recruited to participate in the experiment. Students were approached during economics lectures and were given an information sheet with brief background about the study and requesting their voluntary participation. Students agreeing to participate had to be over 18 years of age, and had to be students at Wits. Since the study design included a test and control group, the experiment was designed to start with a fairly homogenous sample in terms of education, then to randomly assign participants to one of the two groups. To avoid introducing possible biases in responses, the presence of the test and

[4] Note that the existing axiomatizations of ambiguity attitudes within the Anscombe-Aumann (1963) framework (e.g. Schmeidler 1989; Gilboa and Schmeidler 1989; Ghirardato et al. 2004) only allow for violations of the STP but not for violations of the IA whenever an ambiguous decision situation reduces to a risky decision situation.

control groups was not discussed with participants, nor was the exact nature of the experiment. Participants were simply told that this was a study on decision making.

In total, 63 students were recruited, allowing for a minimum of 30 students in each group (31 in the test group and 32 in the control group). The control group had 56% male and 44% female participants, while the test group had 48% male and 52% female participants. In line with Wits' ethical policy on experiments, participants were assured of anonymity in the experiment whereby questionnaire responses were recorded with numbers instead of names.

1.3 Stimuli and statistical learning

We used an actual urn that contained 20 red, 20 yellow, and 60 blue balls. Although the participants did not know the true proportions of the colors in the urn, they were informed that there were 100 balls in total. All respondents answered 30 questions about their preferred choice between two different prospects. By varying the monetary payoffs associated with these prospects (measured in South African Rand whereby R10 $\simeq$ US$ 1)[5], these 30 questions were organized as 15 subsequent choice pairs ($\mathbf{A}, \mathbf{B}; \mathbf{A}', \mathbf{B}'$) exhibiting the payoff structure described in Section 1.1. This design thus allowed for the observation of up to 15 subsequent violations of the STP by any subject through revealed preferences (2).

Following the draw of a ball after each question, respondents in the test group received feedback on the color of the ball drawn from the urn. Respondents from the control group did not receive such feedback. In contrast to the control group, the test group could thus observe statistical information in the form of 30 actual drawings (with replacement) from the urn. Technically speaking, the respondents from the test group were thus exposed to 30 successive iid multivariate Bernoulli trials such that the true proportions of differently colored balls were (supposedly) driving the data generation process.

1.4 Procedure

The respondents sat in front of the computer where they first read participant instructions and then went through the 30 questions of the questionnaire (see the Supplementary Appendix). The first author was present at all times to answer questions and to assist with the random ball selection following each question for the test group.

[5] While we fixed the worst payoff x at zero, the second best payoff y was in the range of R40 to R80, and the best payoff z was in the range of R75 to R150. The Supplementary Appendix lists the questionnnaire containing all 30 questions.

We used randomization of the question pair order to avoid any bias from possible order effects due to, e.g. different magnitudes or ratios from the prospects' possible payoffs. More specifically, the computer programme selected a random starting question pair, and randomized each subsequent question pair. In this way, each respondent would see a unique order of questions, with all respondents seeing all 15 pairs of questions, but in a random order. To avoid bias from the order of presentation within question pairs the labelling of prospects within question pairs was varied.

Once an option had been selected, the test group respondents were allowed to draw a ball from the urn and the computer programme would show the payout based on the color of ball drawn and the prospect selected. The control group received no feedback.

To give incentives for the truthful revelation of preferences, subjects in both groups were told that one of the prospect choices would be selected at random to be paid out in real money (cf., e.g. Starmer and Sugden 1991). The possible payout could be as low as R0 or as high as R150, since these were the minimum and maximum payouts across the range of questions.

At the end of the experiment, i.e. after answering all 30 questions and receiving (the test group) versus not receiving (the control group) statistical information, the respondents from both groups were asked about their estimates for the balls' proportions in the urn.

2 Results

2.1 Statistical learning

We could confirm that statistical learning did indeed happen within the test group. According to standard models of statistical learning, subjective estimates converge to the true proportions of differently colored balls (='objective' probabilities) if the respondents can observe large data samples from multivariate Bernoulli trials (cf., e.g. Viscusi and O'Connor 1984; Viscusi 1985; Chapter 4 and references in Zimper 2013). Answers close to the true proportions of 60% blue, 20% red, and 20% yellow balls in the test group would thus suggest that statistical learning happened within the test group. In contrast, given the absence of any statistical information, the control group was expected to report roughly equal proportions of all colors as an expression of their ignorance.

Table 1: Average estimates for the proportions of differently colored balls

	% Red	% Yellow	% Blue
Control group	33	31	36
Test group	25	21	54

Both patterns are indeed reflected by the reported proportions (Table 1); for example, the average estimate of the test group for the proportion of blue balls was 54% compared to 36% average estimate of the control group. We therefore conclude that exposing the test group to the outcomes of 30 multivariate Bernoulli trials was indeed sufficient to distinguish the test from the control group with respect to statistical learning.

2.2 Revealed choices and STP violations

Table 2 below presents the percentage of choice pairs seen in the control and test group, respectively. Since the question order was randomized to avoid order biases, these question pairs refer to the first, second, etc. pair of questions seen by each respondent. That is, each respondent from the test group had observed 28 subsequent drawings (with replacement) of balls before he/she answered the 15th (=final) question pair.

Table 2: Revealed choices (in %) over rounds of questions

	Control Group				Test Group			
Question pair	$\mathbf{A}, \mathbf{A}'$	$\mathbf{B}, \mathbf{B}'$	$\mathbf{A}, \mathbf{B}'$	$\mathbf{B}, \mathbf{A}'$	$\mathbf{A}, \mathbf{A}'$	$\mathbf{B}, \mathbf{B}'$	$\mathbf{A}, \mathbf{B}'$	$\mathbf{B}, \mathbf{A}'$
1	6	72	16	6	10	58	16	16
2	6	69	19	6	6	48	39	6
3	6	72	9	13	16	48	23	13
4	6	56	19	19	13	42	39	6
5	6	69	19	6	13	45	32	10
6	6	78	16	0	10	35	52	3
7	6	78	13	3	13	42	35	10
8	9	59	25	6	19	45	32	3
9	6	72	22	0	16	39	32	13
10	9	75	16	0	16	45	35	3
11	9	81	9	0	10	42	35	13
12	13	75	9	3	13	45	35	6
13	9	75	16	0	16	45	32	6
14	3	75	22	0	23	29	35	13
15	9	81	6	3	23	55	16	6

Recall that the choice pairs $\mathbf{A}, \mathbf{A}'$ and $\mathbf{B}, \mathbf{B}'$ are consistent with the STP whereas the choice pairs $\mathbf{A}, \mathbf{B}'$ and $\mathbf{B}, \mathbf{A}'$ are not. According to Table 2, the majority of respondents in both groups do not violate the STP; whereby we put down the higher proportion of STP consistent responses in the control group to random preference differences. Among the STP violating responses in both groups, the majority chooses $\mathbf{A}, \mathbf{B}'$. That is, in both groups the majority of our elicited STP violations expresses optimistic rather than pessimistic ambiguity attitudes.

2.3 Measuring the impact of learning

Turn now to our two research questions about whether or not STP violations decrease through learning by mere thought and statistical learning, respectively. Because the control group is exclusively subjected to learning through thought, any decrease in STP violations seen in this group can be solely attributed to learning through mere thought.

We would further assume that learning through thought occurs at the same rate for the test and the control group since both groups are exposed to the same questions in each round. They are also exposed to the same number of such experimental rounds.

While both groups start out in round zero with the same amount (i.e. complete lack) of statistical information, the test group exclusively is exposed to the treatment of statistical learning. Thus, we can isolate the impact of statistical learning by looking at how the differences between STP violations in the test and control groups evolve over rounds.

2.3.1 Overall impact of learning

At first, we look at how the different rounds of learning–thought only for the control and thought as well as statistical learning for the test group–affected the number of STP violations in each group respectively. By the chart in Figure 1, STP violations do not show any clear downward trend for the test group whereas there is a mild downward trend for the control group (albeit with some noise).

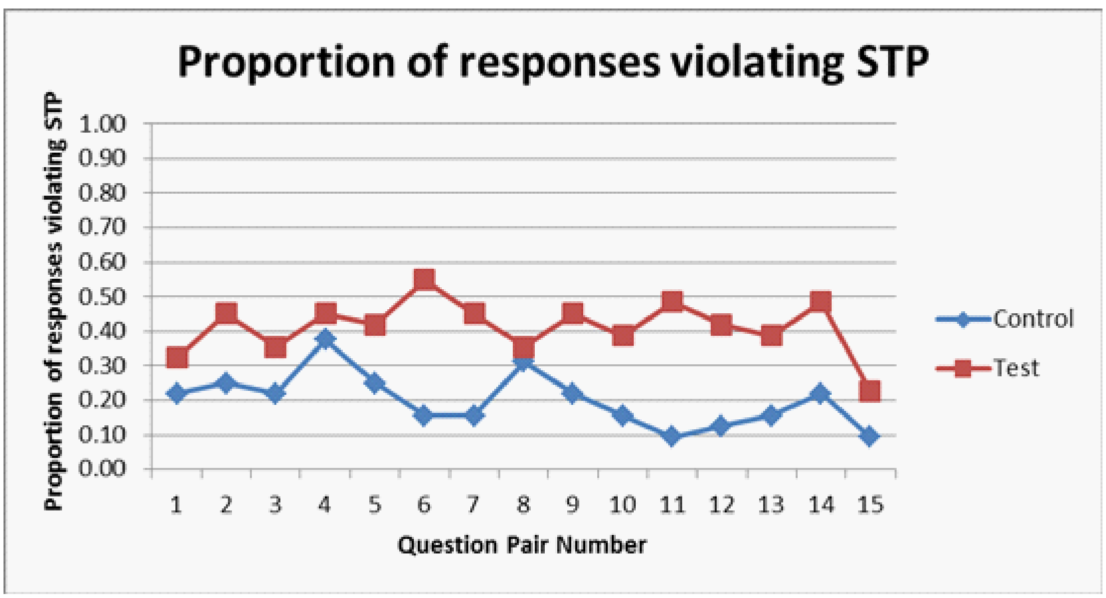

Figure 1: Proportion of responses violating the STP over the course of the experiment for the test and control groups

To obtain a more precise picture, we estimate, at first, the following equation for our group of respondents as a whole:

$$P = \beta_0 + \beta_1 D + \beta_2 D Round + \beta_3 (1 - D) \, Round + \varepsilon \tag{3}$$

where P denotes the proportion (in decimals) of respondents violating the STP, *Round* stands for the number of question pairs answered, and $\varepsilon \sim N(0, \sigma^2)$. We set the dummy variable $D = 1$ for the test and $D = 0$ for the control group so that β_1 measures the initial difference between the test and the control group, β_2 measures the impact on STP violations of an increase in the number of question pairs for the test group, and β_3 measures the same impact for the control group. We estimate both an OLS regression and an iterated reweighted least squares regression (IRLS).[6] The results are reported in Table 3.

Table 3: Estimates of the OLS and IRLS regressions for the proportion of STP violations

	β_0	β_1	β_2	β_3
OLS	0.28*(0.004)	0.14* (0.057)	-0.002(0.004)	-0.01*(0.004)
IRLS	0.27*(0.004)	0.13*(0.057)	0.002(0.004)	-0.009*(0.004)

Notes: standard errors in parentheses
* significant at the 5% level

The results for the OLS and IRLS regressions are very similar. The significant positive coefficient β_1 on the test group dummy shows that there are fewer STP violations in the control than in the test group. For the IRLS regression, in particular, the coefficient implies that the test group is violating the STP by 13 percentage points more than the control group. The estimated β_2 coefficient is very small and insignificant suggesting that the number of experimental rounds for the test group has zero impact on STP violations. In contrast, the negative coefficient β_3, which measures changes in STP violations over rounds for the control group, is significant. In particular, the decrease in STP violations is about 1 percentage point per round for the control group.

In a next step, we conducted a maximum likelihood estimation in the form of logit regression for our individual respondents. The coefficients of a logit regression represent the change in the log of the odds associated with a unit change in the explanatory variable. For ease of interpretation, we also report the marginal effects (ME), which show the increase in the probability, at the means of the covariates, of an individual violating the STP. We conduct both a logit regression on the pooled sample and a fixed effects logit, which allows us to control for idiosyncratic characteristics of the individual that is fixed across rounds, yet might influence his/her probability of violating the STP. We report these results in Table 4.

[6] Compared to the OLS regression, the IRLS regression produces 'robust' estimates in the sense that it accounts for and reduces the influence of extreme observations (observations with large residuals).

Table 4: Logit Estimates

	$\boldsymbol{\beta}_0$	$\boldsymbol{\beta}_1$	$\boldsymbol{\beta}_2$	$\boldsymbol{\beta}_3$
POOLED LOGIT	0.91*(0.227)	0.63*(0.299)	-0.01(0.021)	-0.063*(0.0044)
ME POOLED LOGIT		0.131*	-0.002	-0.01*
FIXED EFFECTS LOGIT			-0.01(0.024)	-0.082*(0.031)
ME FIXED EFFECTS LOGIT			-0.002	-0.02*

Notes: standard errors in parentheses
* significant at the 5% level

These regressions confirm that the number of experimental rounds has a basically zero effect on the probability of an individual in the test group violating the STP. Again, the pooled logit confirms that the probability of an individual in the test group violating the STP is 13 percentage points greater than that of an individual in the control group. We also confirm that the probability of violating the STP decreases over rounds for the control group with the effect being stronger when we control for individual fixed effects. In particular, the fixed effects logit predicts that with every round, the probability of an individual in the control group violating the STP decreases by 2 percentage points.[7]

2.3.2 Isolating the impact of statistical learning

Turn now to the differences between STP violations in the test and the control group measured as the proportion of STP violations in the test group minus the proportion of STP violations in the control group. Figure 2 shows how these differences change over rounds.

[7] Note that there is no $\boldsymbol{\beta}_1$ coefficient for the fixed effects logit since the dummy variable, D, representing whether an individual is in the test or control group, is fixed over rounds and is thus dropped from the regression.

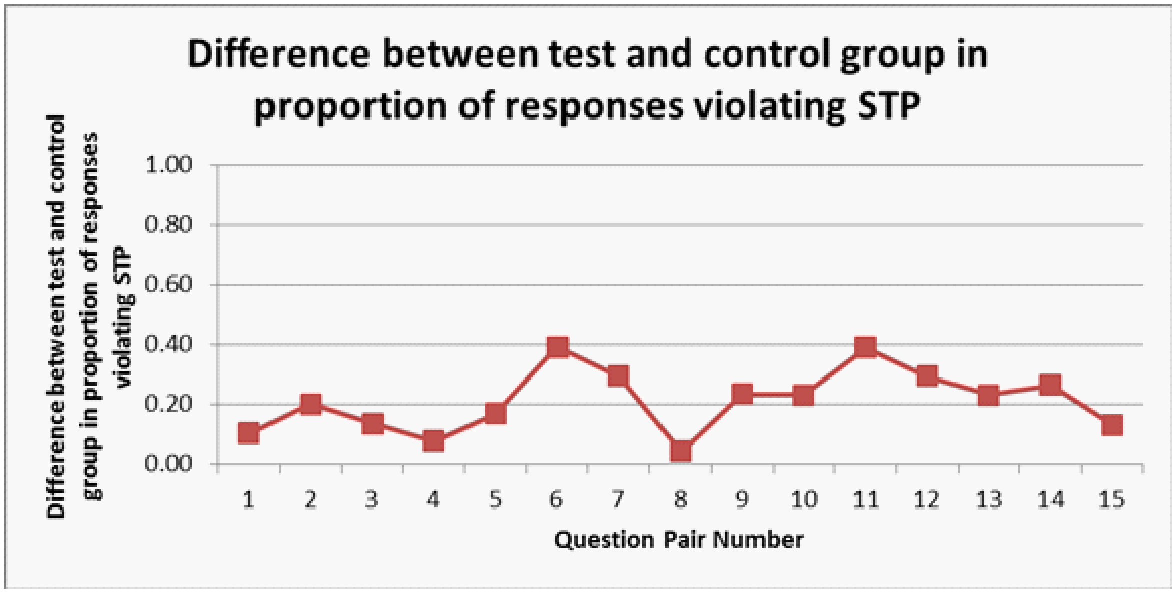

Figure 2: Difference between test and control groups in the proportion of responses violating the STP over the course of the experiment

By Figure 2, there does not appear to be a decrease in this difference, and if there is any change at all, it is a mild increase. To get a more precise picture, we run the following regression

$$DP = \beta_0 + \beta_1 Round + \varepsilon \tag{4}$$

where DP denotes the difference in the proportions (in decimals) of STP violations between Test and Control groups, *Round* stands for the number of question pairs answered, and $\varepsilon \sim N(0, \sigma^2)$. Again, we estimate both an OLS and an IRLS regression. The results are reported in Table 5.

Table 5: Estimates of the OLS and IRLS regressions for the difference in proportion of STP violations

	β_0	β_1
OLS	0.15*(0.056)	0.008 (0.006)
IRLS	0.13 (0.065)	0.009 (0.007)

Notes: standard errors in parentheses
* significant at the 5% level

The β_1 coefficient in both regressions shows a mild but statistically insignificant increase in the difference in the proportion of STP violations between the test and control groups. In particular, the point estimate in the IRLS regression implies that there is an almost 1 percentage point increase per round in the difference in STP violations between test and control groups. The statistical insignificance of this coefficient might in part be due to a relatively large standard error.

2.4 Summary

At the outset of our experiment, the test group exhibits more STP violations than the control group. As the number of experimental rounds increases, the number of STP violations of the control group decreases whereas those of the test group remain approximately constant. As a consequence, the difference in the proportion of STP violations between test and control groups is not decreasing but rather slightly increasing. We ran different types of regressions that all agree on the following two main findings:

1. Learning through mere thought has a small (but significant) positive impact on STP violations in the sene that it leads to a slight decrease of STP violations over the course of the experiment.

2. In contrast, statistical learning has, at best, no impact on STP violations. At worst, it might even be causing STP violations to increase.

3 Discussion

While our conjecture that learning by thought should not lead to an increase in STP violations was confirmed, we were surprised that statistical learning did not lead to such a decrease. With respect to our experimental data both findings appear to be quite robust. Of course, we cannot rule out the possibility that (for whatever reasons) the subjects in our experiment are not representative of a wider population. In particular, observe that the majority of the STP violating choices express optimistic rather than pessimistic ambiguity attitudes. This is in contrast with a large body of experimental literature (cf. Wakker 2010 and references therein) which finds that pessimistic ambiguity attitudes occur more often than optimistic ones. However, even if we restrict attention to the choice pairs $\mathbf{B}, \mathbf{A}'$, which stand for pessimistic ambiguity attitudes, we cannot find a decrease of these STP violating choices through statistical learning. It would be interesting to repeat our (or a similar) experiment with a larger number of subjects who express pessimistic ambiguity attitudes to further validate (or invalidate) our findings.

For two different reasons our experimental design might actually be biased towards a decreasing difference between the proportion of STP violations in the test versus the control group. First, we formally argue in Appendix A3 that the STP violating choice patterns (2) become observationally equivalent to violations of the IA whenever the decision situation under uncertainty has been transformed into a decision situation with known probabilities. Under our maintained assumption that ambiguity should decrease with statistical learning, the test but not the control group's decision situation would have been gradually transformed through statistical learning towards a decision situation with known probabilities. But for a decision situation with known probabilities, van de Kuilen and Wakker (2006) report a decrease in violations of the IA through repeated decision rounds. That is, if statistical learning were to go on for sufficiently many rounds, we would expect a decrease of the choice patterns (2) in the test (and not necessarily in the control) group simply because the test group subjects should decrease their violations of the IA through learning by experience as observed by van de Kuilen and Wakker (2006).

Second, according to our caveat (cf. Appendix A4) we cannot completely rule out the possibility that any decline in STP violating choice patterns (2) for the test group was not exclusively caused by a decrease in STP violations but in addition by the update dynamics of EU consistent decision making for which our experiment did not control. That is, in contrast to the control group, converging statistical learning of EU decision makers in the test group could also have caused a decline in choice patterns (2).

Since our experimental design seems to be biased towards a decrease rather than an increase in the difference of STP violations between the test and the control group, we conclude that the observed absence of such a decrease further confirms the robustness of our main finding.

4 Concluding remarks

This paper has investigated the impact of learning by thought, resp. statistical learning, on violations of Savage's (1954) sure-thing principle (STP). To this purpose we have conducted an experiment which gradually gives more statistical information to the test group in order to supplement this group's learning by thought in the form of repeated exposure to similar choice situations. In contrast, the control group did not receive any statistical information and was thus only subjected to learning by thought.

Our experimental data confirms our first conjecture, according to which learning by thought should not lead to an increase in STP violations. The data even showed a small but significant decrease in STP violations through learning by mere thought. However, the data also rejects our second conjecture, according to which statistical learning should lead to an additional decrease in STP violations. The fact that we cannot find a decrease in STP violations through statistical learning is puzzling to us and, at this point, we do not possess any plausible explanation for this finding.

Appendix

A1. The sure-thing principle

The Savage framework considers a state space S, a set of outcomes X, and a set of acts F which map the state space into the outcome space, i.e. $f : S \to X$ for $f \in F$. Savage (1954) provides a set of axioms over preferences $\succsim$ on the acts in F which gives rise to his famous subjective EU theory. His key axiom for deriving subjective EU theory is the *sure-thing principle* (STP) which formally states that for all acts $f, g, h, h' \in F$ and all events $E \subset S$,

$$f_E h \succsim g_E h \text{ implies } f_E h' \succsim g_E h'. \tag{5}$$

Here $f_E h$ denotes the act that gives the outcomes of act f in event E and the outcomes of act h else, i.e.

$$f_E h(s) = \begin{cases} f(s) & \text{for } s \in E \\ h(s) & \text{for } s \in S \backslash E. \end{cases} \tag{6}$$

To formally describe the choice between the prospects $\mathbf{A}$ and $\mathbf{B}$, resp. $\mathbf{A'}$ and $\mathbf{B'}$, within the Savage framework, let us reinterpret these prospects as the following acts

$$\mathbf{A} = f_{Y \cup B} h, \mathbf{B} = g_{Y \cup B} h; \tag{7}$$
$$\mathbf{A'} = f_{Y \cup B} h', \mathbf{B'} = g_{Y \cup B} h' \tag{8}$$

such that

$$f(s) = \begin{cases} x & \text{for } s \in Y \\ z & \text{for } s \in B \end{cases} \tag{9}$$

as well as $g(s) = y$ for $s \in Y \cup B$ and $h(s) = y$, $h'(s) = x$ for $s \in R$. For example, the revealed choice pair $\mathbf{A}, \mathbf{B'}$, i.e.

$$\mathbf{A} \succ \mathbf{B} \text{ and } \mathbf{B'} \succ \mathbf{A'}, \tag{10}$$

then corresponds to the following revealed preferences over Savage acts

$$f_{Y \cup B} h \succ g_{Y \cup B} h \text{ and } g_{Y \cup B} h' \succ f_{Y \cup B} h', \tag{11}$$

which violate the STP. Similarly, it is straightforward to see that $\mathbf{B}, \mathbf{A'}$ also violates the STP whereas the choice pairs $\mathbf{A}, \mathbf{A'}$ and $\mathbf{B}, \mathbf{B'}$ are consistent with the STP.

A2. Prospect theory

Although the STP violating choice pairs (2) cannot be represented by EU theory, they can be represented by prospect theory. Restricted to the domain of gains, prospect theory gives rise to the following utility representation for preferences $\succsim$ over Savage acts $f, g \in F$,

$$f \succsim g \Leftrightarrow \int_{s \in S}^{C} u\left(f\left(s\right)\right) d\nu\left(s\right) \geq \int_{s \in S}^{C} u\left(g\left(s\right)\right) d\nu\left(s\right) \tag{12}$$

where ν is a unique non-additive (=not necessarily additive) probability measure satisfying $\nu\left(\emptyset\right) = 0$, $\nu\left(S\right) = 1$, and $\nu\left(E\right) \leq \nu\left(E'\right)$ if $E \subset E'$; and the integral in (12) is the Choquet integral.[8]

Focus at first on the choice pair $\mathbf{A}, \mathbf{B}'$. Without loss of generality, set $u\left(z\right) = 1$ and $u\left(x\right) = 0$ so that we obtain for a prospect theory decision maker that

$$\begin{aligned} \mathbf{A} &\succ \mathbf{B} \Leftrightarrow & (14) \\ \int_{s \in S}^{C} u\left(f_{Y \cup B} h\right) dv\left(s\right) &> \int_{s \in S}^{C} u\left(g_{Y \cup B} h\right) dv\left(s\right) \Leftrightarrow & (15) \\ \nu\left(B\right) + u\left(y\right)\left[\nu\left(B \cup R\right) - \nu\left(B\right)\right] &> u\left(y\right) & (16) \end{aligned}$$

as well as

$$\begin{aligned} \mathbf{B}' &\succ \mathbf{A}' \Leftrightarrow & (17) \\ \int_{s \in S}^{C} u\left(g_{Y \cup B} h'\right) dv\left(s\right) &> \int_{s \in S}^{C} u\left(f_{Y \cup B} h'\right) dv\left(s\right) \Leftrightarrow & (18) \\ u\left(y\right)\left[\nu\left(B \cup Y\right)\right] &> \nu\left(B\right). & (19) \end{aligned}$$

Combining inequalities (16) and (19) shows that the utility representation (12) of the revealed choices $\mathbf{A}, \mathbf{B}'$ requires *local concavity* of ν in the following sense[9]

$$\begin{aligned} u\left(y\right)\left[\nu\left(B \cup Y\right)\right] &> u\left(y\right)\left[1 - \left[\nu\left(B \cup R\right) - \nu\left(B\right)\right]\right] \Leftrightarrow & (21) \\ \nu\left(B \cup Y\right) &> \nu\left(S\right) - \left[\nu\left(B \cup R\right) - \nu\left(B\right)\right] \Leftrightarrow & (22) \\ \nu\left(B \cup Y\right) + \nu\left(B \cup R\right) &> \nu\left(\left(B \cup Y\right) \cup \left(B \cup R\right)\right) + \nu\left(\left(B \cup Y\right) \cap \left(B \cup R\right)\right). & (23) \end{aligned}$$

[8] If f only takes on finitely many, say m, values, we have

$$\int_{s \in S}^{C} u\left(f\left(s\right)\right) d\nu\left(s\right) = \sum_{i=1}^{m} u\left(f\left(s_i\right)\right) \cdot \left[\nu\left(E_1 \cup ... \cup E_i\right) - \nu\left(E_1 \cup ... \cup E_{i-1}\right)\right] \tag{13}$$

where $E_1, ..., E_m$ denotes the unique partition of S with $u\left(f\left(s_1\right)\right) > ... > u\left(f\left(s_m\right)\right)$ for $s_i \in E_i$ (cf. Schmeidler 1986).

[9] ν is globally concave (resp. convex) if, and only if, it satisfies for all events E, E'

$$\nu\left(E \cup E'\right) + \nu\left(E \cap E'\right) \leq \text{(resp. } \geq \text{)}\, \nu\left(E\right) + \nu\left(E'\right). \tag{20}$$

An analogous argument shows that the choice pair $\mathbf{B}, \mathbf{A}'$ requires *local convexity* of ν in the following sense

$$\nu(B \cup Y) + \nu(B \cup R) < \nu((B \cup Y) \cup (B \cup R)) + \nu((B \cup Y) \cap (B \cup R)). \tag{24}$$

Concavity, resp. convexity, of the non-additive probability measure ν is typically associated with optimistic, resp. pessimistic, ambiguity attitudes of a prospect theory decision maker (cf. Wakker 2001; Chapter 10.4.3 in Wakker 2010).

A3. Risk and the independence axiom

Consider the limiting case of statistical learning in which the decision maker could observe an infinite number of drawings (with replacement) from the urn, i.e. an infinite amount of data generated by multivariate Bernoulli trials driven by the balls' true proportions. According to standard models of statistical Bayesian learning, this decision maker will then learn, by Doob's (1949) consistency theorem, with certainty the true probabilities of the events R, Y, and B.[10]

Denote this objective probability measure as π^*; e.g. in our experiment we stipulate that $\pi^*(R) = \pi^*(Y) = 0.2$, $\pi^*(B) = 0.6$ in accordance with the balls' true proportions. In this risky decision situation with known probabilities, the prospects $\mathbf{A}$, $\mathbf{B}$, $\mathbf{A}'$, and $\mathbf{B}'$ can be reinterpreted as the following *lotteries* (=objective probability distributions) over the monetary outcomes x, y, z, respectively:

$$\begin{aligned}
\mathbf{A} &= (0.2, 0.2, 0.6), \\
\mathbf{B} &= (0, 1, 0), \\
\mathbf{A}' &= (0.4, 0, 0.6), \\
\mathbf{B}' &= (0.2, 0.8, 0).
\end{aligned}$$

EU theory had been first axiomatized for preferences over lotteries by von Neumann and Morgenstern (1947). Key to their celebrated EU representation theorem is the *independence axiom* (IA) which implies that, for all lotteries L, L', L'' and all $\lambda \in (0, 1)$,

$$L \succsim L' \Leftrightarrow \lambda \cdot L + (1 - \lambda) L'' \succsim \lambda \cdot L' + (1 - \lambda) L''. \tag{25}$$

Focus on the choice pair $\mathbf{A}, \mathbf{B}'$ and observe that it reveals a violation of the IA. To see this denote by δ_k, $k \in \{x, y, z\}$, the degenerate lottery that gives the monetary amount

[10] For a more detailed discussion of the relationship between converging Bayesian learning and Doob's consistency theorem see, e.g. Diaconis and Freedman (1986), Chapter 1 in Gosh and Ramamoorthi (2003), Lijoi, Pruenster and Walker (2004), or Chapter 4 in Zimper (2013). For converging statistical learning in the case of ambiguous priors see Marinacci (2002).

k with probability one. Let $\lambda = \pi^* (R)$ and $(1 - \lambda) \cdot \mu = \pi^* (Y)$ and observe that (25) implies for the choice pair $\mathbf{A}, \mathbf{B}'$ that[11]

$$\begin{aligned} \mathbf{A} \succ \mathbf{B} & \Leftrightarrow & & (26) \\ \lambda \cdot \delta_y + (1 - \lambda) \cdot [\mu \cdot \delta_x + (1 - \mu) \cdot \delta_z] & > & \lambda \cdot \delta_y + (1 - \lambda) \cdot \delta_y \Leftrightarrow & (27) \\ \mu \cdot \delta_x + (1 - \mu) \cdot \delta_z & > & \delta_y . & (28) \end{aligned}$$

as well as

$$\begin{aligned} \mathbf{B}' & \succ & \mathbf{A}' \Leftrightarrow & (29) \\ \lambda \cdot \delta_x + (1 - \lambda) \cdot \delta_y & > & \lambda \cdot \delta_x + (1 - \lambda) \cdot [\mu \cdot \delta_x + (1 - \mu) \cdot \delta_z] \Leftrightarrow & (30) \\ \delta_y & > & \mu \cdot \delta_x + (1 - \mu) \cdot \delta_z . & (31) \end{aligned}$$

Since (28) and (31) constitute a contradiction, the revealed preferences $\mathbf{A}, \mathbf{B}'$ violate the IA.

More generally, in a decision situation with known probabilities the choice pairs (1) are consistent with the IA, whereas the choice pairs (2) violate the IA.

A4. Caveat

Whereas the choice pairs (2) unambiguously reveal a violation of the STP for the control group, the situation is slightly different for the test group whose members observe the drawing of one ball after each choice. To see this consider a decision maker of the test group and denote by I_n the information he has received by observing n drawings. If this decision maker chooses, e.g. $\mathbf{A}$ after n and $\mathbf{B}'$ after $n + 1$ drawings, these revealed choices could be rationalized as EU consistent choices as follows:

$$\begin{aligned} \mathbf{A} \succ \mathbf{B} & \Leftrightarrow & & (32) \\ u(x) \cdot \pi (Y \mid I_n) + u(z) \cdot \pi (B \mid I_n) & > & u(y) \cdot (1 - \pi (R \mid I_n)) & (33) \end{aligned}$$

and

$$\begin{aligned} \mathbf{B}' & \succ & \mathbf{A}' \Leftrightarrow & (34) \\ u(y) \cdot (1 - \pi (R \mid I_{n+1})) & > & u(x) \cdot \pi (Y \mid I_{n+1}) + u(z) \cdot \pi (B \mid I_{n+1}) . & (35) \end{aligned}$$

Whereas the two inequalities (33) and (35) cannot simultaneously hold if $\pi (\cdot \mid I_n)$ and $\pi (\cdot \mid I_{n+1})$ are sufficiently similar probability measures, it is possible that (33) and (35)

[11] We assume that the standard principle of *reduction of compound lotteries* applies (cf., e.g. Fishburn 1988).

are satisfied for appropriately chosen values of $\pi(\cdot \mid I_n)$ and $\pi(\cdot \mid I_{n+1})$. In that case, the choices $\mathbf{A}, \mathbf{B}'$ would not reveal a violation of the STP but rather the different perception of the urn's uncertainty by an EU decision maker before versus after he updates his additive belief on the $n+1$th observation.

We could have easily avoided this ambiguity in the interpretation of choice pairs (2) for the test group, if we had allowed for statistical learning not within but only after each question pair was answered. However, when designing the experiment, our concern was to give the subjects no hint that they were actually answering 15 well-structured question pairs rather than 30 similar questions. Although we feared that a detection of this question pair structure might eventually result in some answering bias, we assumed that any difference between $\pi(\cdot \mid I_n)$ and $\pi(\cdot \mid I_{n+1})$ would be negligibly small (in the sense of: How great can the impact of a single observation be?).

In the subsequent interpretation of the data, we therefore assume that $\pi(\cdot \mid I_n)$ and $\pi(\cdot \mid I_{n+1})$ are indeed sufficiently close so that choice pairs (2) cannot be explained by EU consistent decision making but rather by revealed violations of the STP. If this assumption was violated, however, we would observe, by Doob's (1949) consistency theorem, that $\pi(\cdot \mid I_n)$ and $\pi(\cdot \mid I_{n+1})$ become more and more similar with increasing n. Applied to the experiment, this means that although revealed choices (2) in the first rounds of the experiment (small n) might be EU consistent, revealed choices (2) in the later rounds of the experiment (large n) would rather indicate a violation of the STP. Consequently, an EU consistent decision maker who expressed choice pairs (2) in the beginning of the experiment would eventually switch to choice pairs (1) in the later rounds of the experiment when his conditional probability measures converge through statistical learning in accordance with Doob's (1949) consistency theorem.

To summarize the caveat: If our maintained assumption that choice pairs (2) always reveal violations of the STP for the test group was not correct, we might observe a decline in the number of choice pairs (2) which is not caused by a decline of violations of the STP.

References

Abdellaoui, M., Baillon, A., Placido, L., and Wakker, P. P. (2011) "The Rich Domain of Uncertainty: Source Functions and their Experimental Implementation," *American Economic Review* **101**, 695–723.

Allais, M. (1979) "The Foundation of a Positive Theory of Choice Involving Risk and a Criticism of the Postulates and Axioms of the American School". Part II in Allais, M., and O. Hagen [eds.], *Expected Utility Hypotheses and the Allais Paradox*, D. Reidel: Dordrecht.

Anscombe, F. J., and R. J. Aumann (1963) "A Definition of Subjective Probability," *Annals of Mathematical Statistics* **34**, 199-205.

Diaconis, P., and D. Freedman (1986) "On the Consistency of Bayes Estimates," *Annals of Statistics* **14**, 1-26.

Doob, J. L. (1949) "Application of the Theory of Martingales," In: *Le Calcul des Probabilites et ses Applications*, Colloques Internationaux du Centre National de la Recherche Scientifique 13, pp. 23–27. Paris: CNRS.

Ellsberg, D. (1961) "Risk, Ambiguity and the Savage Axioms," *Quarterly Journal of Economics* **75**, 643-669.

Epstein L. G., and M. Schneider (2007) "Learning under Ambiguity," *Review of Economic Studies* **74**, 1275–1303.

Fishburn, P. C. (1988) *Nonlinear Preferences and Utility Theory under Risk*, The John Hopkins University Press: Baltimore, London.

Ghirardato, P. and M. Marinacci (2002), "Ambiguity Made Precise: A Comparative Foundation," *Journal of Economic Theory* **102**, 251-289.

Ghirardato, P., Maccheroni, F., and M. Marinacci, M. (2004) "Differentiating Ambiguity and Ambiguity Attitude," *Journal of Economic Theory* **118**, 133–173.

Gilboa, I. (1987) "Expected Utility with Purely Subjective Non-Additive Probabilities," *Journal of Mathematical Economics* **16**, 65-88.

Gilboa, I., and D. Schmeidler (1989) "Maxmin Expected Utility with Non-Unique Priors," *Journal of Mathematical Economics* **18**, 141-153.

Gosh, J. K., and R. V. Ramamoorthi (2003) *Bayesian Nonparametrics*, Springer Verlag Inc.: New York.

Jaffray, J-Y. (1994) "Dynamic Decision Making with Belief Functions," In: R. R. Yager, M. Fedrizzi, and J. Kacprzyk [eds.], *Advances in the Dempster-Shafer Theory of Evidence*, pp.331–352. Wiley, New York.

Lijoi, A., Pruenster, I., and S.O. Walker (2004), "Extending Doob's Consistency Theorem to Nonparametric Densities," *Bernoulli*, **10**, 651-663.

Machina, M. J. (2004) "Almost-Objective Uncertainty," *Economic Theory* **24**, 1–54.

Marinacci, M. (2002) "Learning from Ambiguous Urns," *Statistical Papers* **43**, 145-151.

Savage, L. J. (1954) *The Foundations of Statistics*, John Wiley and & Sons, Inc.: New York, London, Sydney.

Schmeidler, D. (1986) "Integral Representation without Additivity," *Proceedings of the American Mathematical Society* **97**, 255-261.

Schmeidler, D. (1989) "Subjective Probability and Expected Utility without Additivity," *Econometrica* **57**, 571-587.

Schmidt, U. (2004) "Alternatives to Expected Utility: Some Formal Theories," in: P.J. Hammond, S. Barberá, and C. Seidl [eds.], *Handbook of Utility Theory* Vol. II, Kluwer, Boston, Chapter 15.

Starmer, C. (2000) "Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk," *Journal of Economic Literature* **38**, 332-382.

Starmer, C., and R. Sugden (1991) "Does the Random-Lottery Incentive System Elicit True Preferences? An Experimental Investigation," *American Economic Review* **81**, 971–978.

Sugden, R. (2004) "Alternatives to Expected Utility: Foundation," in: P.J. Hammond, S. Barberá, and C. Seidl [eds.], *Handbook of Utility Theory* Vol. II, Kluwer, Boston, Chapter 14.

Tversky, A., and D. Kahneman (1992) "Advances in Prospect Theory: Cumulative Representations of Uncertainty," *Journal of Risk and Uncertainty* **5**, 297-323.

Viscusi, W. K. (1985) "A Bayesian Perspective on Biases in Risk Perception," *Economics Letters* **17**, 59-62.

Viscusi, W. K. and C. J. O'Connor (1984) "Adaptive Responses to Chemical Labeling: Are Workers Bayesian Decision Makers?," *American Economic Review* **74**, 942-956.

van de Kuilen, G., and P. P. Wakker (2006) "Learning in the Allais Paradox," *Journal of Risk and Uncertainty* **33**, 155-164.

von Neuman, J., and O. Morgenstern (1947) *Theory of Games and Economic Behavior*, Princeton University Press, 2nd edn., Princeton.

Wakker, P. P (2001) "Testing and Characterizing Properties of Nonadditive Measures through Violations of the Sure-Thing Principle," *Econometrica* **69**, 1039-1059.

Wakker, P. P (2010), *Prospect Theory for Risk and Ambiguity*, Cambridge University Press.

Wakker, P. P, and A. Tversky (1993) "An Axiomatization of Cumulative Prospect Theory," *Journal of Risk and Uncertainty* **7**, 147-167.

Wu, G., and R. Gonzalez (1996) "Curvature of the Probability Weighting Function," *Management Science* **42**, 1676-1690.

Wu, G., and R. Gonzalez (1999) "Nonlinear Decison Weighst in Choice under Uncertainty," *Management Science* **45**, 74-85.

Zimper, A. (2011) "Do Bayesians Learn their Way out of Ambiguity?," *Decision Analysis* **8**, 269-285.

Zimper, A. (2013) "The Emergence of "Fifty-Fifty" Probability Judgments through Bayesian Updating under Ambiguity," *Fuzzy Sets and Systems* **223**, 72-88.

Zimper, A., and A. Ludwig (2009) "On Attitude Polarization under Bayesian Learning with Non-additive Beliefs," *Journal of Risk and Uncertainty* **39**, 181-212.