Ellsberg Meets Nash: the Ellsberg Task as a Game

Adam Dominiak*
Department of Economics
Virginia Tech

Peter Duersch†
Department of Economics
University of Heidelberg

This Version: October 14, 2015

Abstract

In his PhD thesis, Ellsberg formulated strong experimental conditions for his proposed tests of subjective expected utility theory. Subjects should have no reason to consider the motives of the urn filler. Standard incentivized experiments do not meet these conditions. Instead of a one-person decision problem, the task can be perceived as a two-player game. One player chooses among the bets. The second player determines the distribution of balls. The Nash equilibrium predictions depend on the payoff of the second player, yielding a zero-sum or a coordination game. Implementing both situations experimentally does not support the ambiguity averse preferences interpretation of the task, which predicts no differences across treatments. To the contrary, subjects' decisions depend on the motives of the urn filler.

Keywords: Ellsberg task, experiment, zero-sum game, coordination game, ambiguity, uncertainty averse preferences

JEL-Codes: C91, D81

---

*dominiak@vt.edu, Virginia Polytechnic and State University, Department of Economics, 3122 Pamplin Hall (0316), 880 West Campus Drive, Blacksburg, VA 24061, USA.

†peter.duersch@awi.uni-heidelberg.de, University of Heidelberg, Department of Economics, Bergheimer Str. 58, D-69115 Heidelberg, Germany.

We are grateful to Colin Camerer, Soo Hong Chew, Jürgen Eichberger, Hans Haller, Jörg Oechssler, and Stefan Trautmann for their valuable comments and discussions. We wish to thank to the participants of ESA Conference in Tucson, the 16th Conference on the Foundations of Risk and Utility (FUR) in Rotterdam, the 5th ICES Conference on Behavioral and Experimental Economics at George Mason University, the Risk, Uncertainty and Decision (RUD) Conference at Bocconi University and the workshops at Hebrew University of Jerusalem and at University of Mannheim for their comments.

1 Introduction

In 1961, Daniel Ellsberg proposed a class of choice problems designed to test the subjective expected utility (SEU) hypothesis formalized by Savage (1954). Ellsberg conjectured that subjects might exhibit choices that are incompatible with Savage's theory. Later experiments implemented Ellsberg's idea and largely reported behavior in line with his predictions. In this paper, we challenge the prevalent interpretation of the Ellsberg choices observed in non-hypothetical experiments. Under a very natural assumption about the motivation and utility function of the experimenter, the Ellsberg behavior is supported by the Nash equilibrium in a game between the subject and the experimenter. We run an experiment testing whether subjects view the Ellsberg task as a game versus the standard view of a one-person decision problem. The data supports the interpretation as a game.

The key element of Ellsberg's thought-experiments is a sealed urn filled with colored balls. The so-called Ellsberg urn is designed to present ambiguity, a choice situation where some states of nature do not have an obvious probability assignment. In the 3-color version, the urn is placed in front of subjects who are informed that 30 balls are of red color and 60 balls are somehow divided between black and yellow balls without any information on the composition. Based on that information, subjects are confronted with two choice problems, together called the Ellsberg task. In the first choice problem, subjects are asked to choose between the bet yielding \$1 if the ball drawn is red and \$0 otherwise; and the bet yielding \$1 if the ball drawn is black and \$0 otherwise. The second choice is between the bet paying \$1 if the drawn ball is not-black and \$0 otherwise; and the bet that paying \$1 if the drawn ball is not-red and \$0 otherwise. After choices are made, two independent draws with replacements are conducted. Subjects are paid according to their choices and the experiment ends.[1]

[1]Ellsberg proposed two versions of his thought-experiments: one with 2 colors and 2 urns and one with 3 colors and 1 urn. The paper focuses on the 3-color version. The theoretical results carry over to the 2-color experiment, see Appendix 7.

When certain "*experimental conditions*" - as Ellsberg (2001, p.137) in his doctoral thesis called them - are satisfied, the choices made can be used to verify consistency with Savage's key axiom, the Sure-Thing Principle. Ellsberg carefully outlined the hypothetical situation in which his experiments are to take place: "*You are not told by whom, or by what process, or for what purpose the composition of Urn II was determined (except that it was* **not** *determined, or revealed, to the observer who is offering you gambles). [...] In particular, you might try to imagine circumstances that would convince you that the person offering you gambles* **knows no more about the contents of Urn II, or how they are determined, than you do**, *so that his offers convey to you no "information" on those contents*" (Ellsberg, 2001, p.131-132, original emphasis).

Typically, the Ellsberg decision problem is formulated in a formal framework with a state space representing the outcomes of a random draw from the urn. The objects of choice (i.e., the bets) are depicted as mappings from states to the monetary payoffs. Ellsberg showed that betting on red (resp. black) in the first choice and betting on non-red (resp. non-black) in the second choice can not be rationalized by any subjective probability distribution. The systematic preference to *bet on the known* probability colors (*bet-on-K*) has been interpreted as *ambiguity aversion.* Conversely, the preference to *bet on the unknown* probability colors (*bet-on-U*) has been labeled *ambiguity loving.* Betting otherwise (*bet-on-C*) is compatible with probabilistic beliefs and having a belief that one color, either black or yellow, is more likely. This behavior has been called *ambiguity neutral.*

The publication of Ellsberg's thought-experiments paved the way for experimental and theoretical literature on decision making in the presence of ambiguity. A growing number of experimental studies, summarized by Camerer and Weber (1992), implemented Ellsberg's choice problems. This experimental data was interpreted as evidence in favor of ambiguity-sensitive behavior, despite the experiments not adhering to the conditions laid out by Ellsberg.[2] In 1961, there was no

[2]We are not aware of any incentivized study for which the above quotation could be said to be completely true. It is clearly violated in the standard case, when the urn is directly filled by

sound theory that could accommodate his "paradoxical" behavior. It took close to 30 years until the first axiomatically founded model explaining Ellsberg's choices could be established.[3]

The first prominent approach to ambiguity built on the idea that subjective beliefs are represented by a set of probability distributions (priors). In their path-breaking contribution, Gilboa and Schmeidler (1989, p.142) suggested the following explanation of the Ellsbergian behavior: *[...] the subject has too little information to form a prior. Hence (s)he considers a set of priors as possible. Being uncertainty averse, (s)he takes into account the minimal expected utility (over all priors in the set) while evaluating a bet.* The resulting model, known as Maxmin Expected Utility (MEU), is an archetype of ambiguity averse preferences. The model has been extended in many directions leading to the general family of uncertainty averse preferences (UAP) derived by Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio (2011) which includes the multiple player preferences of Hansen and Sargent (2001) as well as the variational preferences of Maccheroni, Marinacci, and Rustichini (2006).

To draw conclusions about ambiguity-sensitive behavior from experiments, Ellsberg's experimental conditions need to be satisfied. However, these conditions, and the formal framework used, translate to two implicit, but questionable, assumptions: First, it is taken for granted that subjects facing the Ellsberg task act in "isolation". That is, subjects act without taking into account the potential incentives - the utility function - of whoever created the distribution in the Ellsberg urn. The motives of the person filling the urn are entirely neglected. Second, it is assumed that subjects evaluate each choice problem separately (i.e., independent from each other). This kind of "separation-principle" is implied by the fact that subjects who choose according to the Ellsberg paradox rationalize their choices by distinct probability distributions; one prior for each choice problem.

---

the experimenters.

[3]For comprehensive surveys on the ambiguity literature see Gilboa and Marinacci (2013); Karni, Maccheroni, and Marinacci (2014); Machina and Siniscalchi (2014).

Focus for the moment on the first assumption. The question arises whether subjects do indeed view the Ellsberg task as a non-interactive decision problem. This is far from being obvious (and should be tested).[4] In the classic experiments used to verify Ellsberg's predictions, an entity filling the urn always exists in the form of the experimenter and his or her motivations can be considered by the subjects.[5] Modeling this situation does not require a new model; there exists an extremely powerful concept that is concerned exactly with the question of how to react to the motivations, the payoffs, of other people: the Nash equilibrium. It will be shown that including the motives of the urn filler accounts for Ellsberg's predictions. Moreover, taking into account the distinct motivations impacts the predictions in a way that can be tested experimentally.

Based on the 3-color experiment, *3-color Ellsberg game* refers to the two-player game where one player fills the Ellsberg urn with balls (adhering to the fixed and non-fixed proportions as outlaid by Ellsberg) and the other player chooses the Ellsberg bets, as described above. Denote by D the player facing the Ellsberg task and by F the player filling the urn with balls. That is, player D's strategy set consists

[4] Hsu, Bhatt, Adolphs, Tranel, and Camerer (2005) conduct a fMRI study to compare neural activations during decisions that involve risk or ambiguity. They find activity differences in the amygdala, the orbitofrontal cortex (OFC) and the dorsomedial prefrontal cortex. In another fMRI study Grèzes, Frith, and Passingham (2004) measure the neural activity when subjects evaluate being deceived by others (relative to evaluating other's beliefs). Among other regions, they find activation in the OFC and the amygdala. Possibly, subjects presented with an ambigious situation try to judge the intentions of another person.

[5] In some experiments the composition of the urn is not determined directly by the experimenter. Instead, some external random generator is used, e.g. a mechanical device or future stock volatility (see Hey, Lotito, and Maffioletti, 2010; Abdellaoui, Baillon, Placido, and Wakker, 2011). However, as long as subjects are paid by the experimenter, subjects may feel that the experimenter, having chosen the random mechanism, may have superior knowledge about the outcome of the random mechanism. Kadane (1992) points out that Ellsberg behavior can be explained if subjects suspect that the experimenter cheats by using procedures and a random mechanism for determining the payoffs which are different from the ones presented to the subjects.

of the four possible combinations of bets, while player F's strategy set consists of all admissible urn compositions. Player D's utility function is given by the Ellsberg bet rewards. Player F's utility depends on the payoff of the Ellsberg bets as well. Modeling different motives of player F, it can be positively or negatively correlated with the payoff of player D, leading to games where D plays *with* or *against* F.

It is noteworthy that outside motivations feature also in the ambiguity models. However, the motivations are inherent in the functional representing the UAP. It is the *min*-operator that implements a negative view of the outside world; a "malevolent Nature". That is, Nature is viewed as being able to influence the occurrence of events to the decision makers' disadvantage by choosing a probability distribution with the objective of minimizing his expected utility. For this reason, UAP are sometimes said to admit *"a game against Nature interpretation, where decision makers view themselves as playing a zero-sum game against (a malevolent) Nature"* (Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio, 2011, p. 1299). The decision maker's strategy is choosing a bet while Nature's strategy is choosing a probability distribution. However, this negative view is a primitive of the model and a feature of the decision maker's preferences, but not a variable of the situation. As such, the UAP models can not accommodate changes in motivation without changing the preference parameters.

Consider now to the second assumption. According to the "separation-principle", an uncertainty averse decision maker may use different probabilities when evaluating the two choice problems of the Ellsberg task. Correspondingly, in the ambiguity models, subjects are viewed as playing two different "games against a malevolent Nature"; a separate one referring to each choice in the 3-color Ellsberg problem. This is not the case in the Nash equilibrium setup where only one game is played.

While Ellsberg and the UAP models use the possible outcomes of draws from the urn as the domain of subjective beliefs, the Nash equilibrium considers actions of the opponent player. As such, player D's beliefs are defined not over the possible colors of the drawn ball, but over the content of the urn before the draw. Considering an opponent and his strategies furthers this form of beliefs. Imagine the possible thought process of a subject in the actual experiment as it takes place in a laboratory: choices are paid; someone pays the subjects money; everyone prefers to pay less to paying more; changing the contents of the urn can, depending on the subjects choices, result in paying less; so, how could the experimenter have filled the urn to pay out less? While looking at the strategies of the urn filler, however, the subject has already taken the step to consider the two Ellsberg choices as a part of one, bigger, problem: how to obtain the highest payoff as a subject, given the action of the urn filler. Therefore, the experimental setup pushes subjects to evaluate the problem differently from what is assumed in the UAP models: because there is interaction with a second person, the choices are integrated into one decision problem instead of treating them as two separate problems. The action of the urn filler becomes the focus.

Since the payoff of any urn filler is not incorporated into the UAP functional, changing this payoff does not directly affect the prediction of the models. On the other hand, Nash equilibria can change when the payoff of one player is altered. In the experimental treatments, payoffs of the urn fillers are varied. This creates different Nash predictions. Note that the information about the urn that decision makers in all treatments receive is identical. In each case there is the same number of known and unknown balls in the urn which is physically constructed in the same way. Consequently, the predictions based on the UAP models remain unchanged across treatments, unless the functional changes.

To vary the urn fillers' payoff, in two of our treatments the urn is not filled by the experimenter, but by one group of subjects. In one treatment *(malevolent)*, these subjects have monetary incentives which are opposite to those of the second group of subjects making the Ellsberg decisions. In this treatment, each pair of urn filler and decision maker face a zero-sum game in terms of payoffs. As is shown in Section 2, the pure Nash equilibrium of this malevolent 3-color Ellsberg game predicts that decision makers *bet-on-K*.[6] In a second treatment *(benevolent)*, the

[6] Interestingly, the Nash equilibrium strategies of the decision maker depend on the risk attitude of the urn filler. While playing *bet-on-K* is a Nash equilibrium strategy under any risk attitude

monetary incentives of the urn filler are perfectly aligned with the decision maker: they play a coordination game. This changes the Nash prediction by introducing two additional coordination equilibria in pure strategies: one for betting on black and one for betting on yellow.

As the main result, the measured betting behavior of decision makers differs between malevolent and benevolent, in line with the Nash prediction. About half of the subjects in the malevolent treatment choose *bet-on-K*. When facing the benevolent urn filler, less than a quarter of subjects decides to *bet-on-K*. Instead, most choose *bet-on-C* strategies. These are the coordination Nash equilibria strategies of the benevolent 3-color Ellsberg game: when the payoffs make coordination with the filler of the urn possible, subjects attempt coordination. When the payoffs of urn filler and decision maker are opposed, the coordination Nash equilibria disappear and the associated strategies are observed less often. This suggests that at least a part of subjects consider their task to be an interactive and integrated one, instead of the non-interactive, separated tasks assumed in the UAP framework. It is important to note that, in the integrated view of the Ellsberg task, subjects need not violate SEU by choosing *bet-on-K*. A third treatment *(standard)* is used to relate the findings to the typical implementation of the Ellsberg urn in experiments. Here, the urn is filled by the experimenter and all subjects are decision makers. Betting behavior differs significantly between the standard and the benevolent treatment, but not between the standard and the malevolent treatment.

This paper is organized as follows: in Section 2, the 3-color Ellsberg experiment is described and analyzed as a game between decision maker and urn filler. Nash equilibria are derived for the malevolent and benevolent treatment. In Section, 3 the alternative prediction by uncertainty averse preference models is presented. Section 4 compares the two approaches theoretically. Section 5 describes the experimental design and results. Section 6 concludes the paper.

---

of the urn filler, one additional Nash equilibrium exists only when the urn filler is risk neutral. In this equilibrium, the decision maker chooses *bet-on-U*.

2 Ellsberg Task and Nash Behavior.

In this section, the classical 3-color experiment of Ellsberg (1961) is interpreted as an interactive decision problem. We introduce the notion of a 3-color Ellsberg game between a subject and an urn filler and examine Nash equilibrium behavior.

The 3-color Ellsberg game consists of two players, indexed by D and F. Both players' decisions involve an urn containing 30 balls, 10 of which are known to be red (R) and 20 of which are somehow divided between black (B) and yellow (Y).[7] Player D represents the decision maker (subject) who faces the 3-color Ellsberg task. The task is to make two choices, each between two bets. The bets are presented in Table 1. Each bet pays off \$1 or \$0, depending on the color of a random and independent draw from the urn.

Table 1: Bets in the 3-color Ellsberg experiment.

		Red	*Black*	*Yellow*
Choice I	f_1	\$1	\$0	\$0
	f_2	\$0	\$1	\$0
Choice II	f_3	\$1	\$0	\$1
	f_4	\$0	\$1	\$1

The two choices made in the 3-color Ellsberg experiment - taken together - can be seen as a (pure) strategy for player D. Four possible choice combinations exist. Therefore, the strategy set of player D is $S_D = \{(f_1, f_3); (f_1, f_4); (f_2, f_3); (f_2, f_4)\}$. The strategy of choosing f_1 in problem I and f_4 in problem II is labeled *bet-on-K*, where K stands for the known probability events R and $B \cup Y$. Conversely, the strategy choosing f_2 in problem I and f_3 in II is called *bet-on-U* where U denotes

[7] For ease of exposition, throughout the paper, the traditional Ellsberg colors, red, black and yellow, and the terms "urn" and "balls" are used. The actual experiment and instructions used yellow, blue and green marbles in cotton bags instead.

the unknown probability events B and $R \cup Y$. The choices of (f_1, f_3) are called *bet-on-Y* and (f_2, f_4) *bet-on-B* with Y and B denoting the yellow and black colors, respectively. Betting on one of the two colors - yellow or black - by using either *bet-on-Y* or *bet-on-B*, is summarized as *bet-on-C*.

The task of player F is to determine the composition of the Ellsberg urn by filling the urn with black and yellow balls. Altogether, there are twenty-one possible distributions. For $n = 0, \ldots, 20$, denote by b_n the number of black balls. If player F decides for b_n black balls, the urn contains 10 red, b_n black and $20 - b_n$ yellow balls. The set of all the possible distributions is denoted by $S_F = \{b_0, b_1, \ldots, b_{20}\}$ and represents the set of (pure) strategies of player F.

Our approach distinguishes two different types of the 3-color Ellsberg game, depending upon the correlation of payoffs of player F with the payoffs of player D. The game where player F's payoffs are perfectly negatively correlated is referred to as the *malevolent* game. The game is called *benevolent* when D's and F's payoffs are perfectly positively correlated.

The malevolent 3-color Ellsberg game is zero-sum in payoffs. That is, whenever the color on which player D decides to bet matches the color of a randomly drawn ball, player D wins \$1 while player F loses \$1. Throughout the analysis, the utility function over monetary payoffs is assumed to be strictly increasing and normalized to 0 for the payoff of \$0 and 1 for the payoff of \$1.[8] Since two independent random draws are carried out, each strategy of player F induces a lottery over payoffs of \$2, \$1 and \$0. For the analysis of Nash behavior, it is assumed that players' preferences are standard expected utility among (risky) lotteries. The expected payoffs are calculated in the following way. Suppose that player D chooses (f_1, f_4) (i.e., *bet-on-K*) and player F decides for the uniform distribution b_{10}. Given this strategy combination, player D gets \$2 with a probability of $\frac{2}{9}$, which is the likelihood of drawing red in the first draw times the probability of drawing black or yellow in the second draw, and \$1 with a probability of $\frac{5}{9}$, which is the probability of drawing

[8]In the actual experiment, the payoffs are 0 EUR and 4 EUR instead of \$0 and \$1.

red in the first and second draws plus the probability of drawing not-red in both draws. Given the uniform distribution of player F, player D's expected payoff from playing (f_1, f_4) amounts to $u(2)\frac{2}{9} + u(1)\frac{5}{9}$. Given D's strategy, player F loses in expectation $u(-2)\frac{2}{9} + u(-1)\frac{5}{9}$ by choosing the uniform distribution. The expected utilities for player D's strategies, given that player F fills the urn with b_n black balls, are

$$\begin{aligned}
EU\big[(f_1, f_3) \mid b_n\big] &= u(2)\tfrac{30-b_n}{90} + u(1)\tfrac{60-b_n}{90}, \\
EU\big[(f_1, f_4) \mid b_n\big] &= u(2)\tfrac{20}{90} + u(1)\tfrac{50}{90}, \\
EU\big[(f_2, f_3) \mid b_n\big] &= u(2)\tfrac{b_n(30-b_n)}{900} + u(1)\tfrac{(b_n)^2+(30-b_n)^2}{900}, \\
EU\big[(f_2, f_4) \mid b_n\big] &= u(2)\tfrac{2b_n}{90} + u(1)\tfrac{60-b_n}{90},
\end{aligned} \tag{1}$$

where $n = 0, \ldots, 20$.

Let us first assume that both players are risk neutral. In this case, the malevolent 3-color Ellsberg game is a zero-sum game. Table 2 represents the expected payoffs associated with each strategy combination. In the malevolent Ellsberg

Table 2: Expected payoffs in the malevolent 3-color Ellsberg game.

	bet-on-Y	***bet-on-K***	***bet-on-U***	*bet-on-B*
b_{20}	$-\frac{20}{30}, \frac{20}{30}$	$-1, 1$	$-1, 1$	$-\frac{40}{30}, \frac{40}{30}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_{15}	$-\frac{25}{30}, \frac{25}{30}$	$-1, 1$	$-1, 1$	$-\frac{35}{30}, \frac{35}{30}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\boldsymbol{b_{10}}$	$-\frac{30}{30}, \frac{30}{30}$	$\boldsymbol{-1, 1}$	$\boldsymbol{-1, 1}$	$-\frac{30}{30}, \frac{30}{30}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_5	$-\frac{35}{30}, \frac{35}{30}$	$-1, 1$	$-1, 1$	$-\frac{25}{30}, \frac{25}{30}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_0	$-\frac{40}{30}, \frac{40}{30}$	$-1, 1$	$-1, 1$	$-\frac{20}{30}, \frac{20}{30}$

game, there are two Nash equilibria in pure strategies: $NE_1 = \big\{\textit{bet-on-K}; b_{10}\big\}$ and $NE_2 = \big\{\textit{bet-on-U}; b_{10}\big\}$. In both equilibria, player F fills the Ellsberg urn with an equal number of black and yellow balls. As a best response to this, player D's chooses *bet-on-K* in the first Nash equilibrium. That is, he bets on the events with known probabilities. In the second equilibrium, player D bets on the events with unknown probabilities and chooses the strategy *bet-on-U*. Both equilibria are payoff equivalent for the players. Under the game-theoretic interpretation of the Ellsberg experiment and risk neutrality, Nash behavior predicts *bet-on-K* and *bet-on-U* together with the uniform composition of the urn.[9]

In the benevolent 3-color Ellsberg game, the payoffs of player D and F are perfectly positively correlated. Both players win \$1 whenever the color of a randomly drawn ball matches the chosen color(s) of player D and \$0 when the colors do not match. Under risk neutrality of both players, the expected payoffs in the benevolent Ellsberg game are depicted in Table 3. Unlike the malevolent game, this is not a zero-sum game. The set of Nash-Equilibria is larger in the benevolent game. Again, $NE_1 = \{\textit{bet-on-K}; b_{10}\}$ and $NE_2 = \{\textit{bet-on-U}; b_{10}\}$ constitute pure Nash equilibria. However, there are two additional Nash equilibria in pure strategies:

Table 3: Expected payoffs in the benevolent 3-color Ellsberg game.

	bet-on-Y	***bet-on-K***	***bet-on-U***	***bet-on-B***
$\boldsymbol{b_{20}}$	$\frac{20}{30}, \frac{20}{30}$	$1, 1$	$1, 1$	$\mathbf{\frac{40}{30}, \frac{40}{30}}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_{15}	$\frac{25}{30}, \frac{25}{30}$	$1, 1$	$1, 1$	$\frac{35}{30}, \frac{35}{30}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\boldsymbol{b_{10}}$	$\frac{30}{30}, \frac{30}{30}$	$\mathbf{1, 1}$	$\mathbf{1, 1}$	$\frac{30}{30}, \frac{30}{30}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_5	$\frac{35}{30}, \frac{35}{30}$	$1, 1$	$1, 1$	$\frac{25}{30}, \frac{25}{30}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\boldsymbol{b_0}$	$\mathbf{\frac{40}{30}, \frac{40}{30}}$	$1, 1$	$1, 1$	$\frac{20}{30}, \frac{20}{30}$

[9] In the analysis, we concentrate on pure Nash equilibia. For sake of completeness, mixed behavior is analyzed in Appendix 9.

$NE_3 = \{bet\text{-}on\text{-}Y; b_0\}$ and $NE_4 = \{bet\text{-}on\text{-}B; b_{20}\}$. The additional Nash equilibria predict coordination behavior. Both players coordinate on one color and obtain the highest possible joint payoff.

It is also worth mentioning that the equilibrium behavior with coordination on the black or yellow color strongly Pareto dominates the behavior in NE_1 and NE_2 where player D chooses *bet-on-K* and *bet-on-U*, respectively. Furthermore, equilibria NE_1 and NE_2 are also risk dominated by the two pure coordination equilibria.

So far, the analysis has been conducted under the assumption of risk neutral players.[10] Now, suppose that players' utility functions are nonlinear. Note that the malevolent game need not be zero-sum in utilities any longer. Table 4 represents player D's expected utilities for each strategy combination.

Table 4: Expected payoffs of player D - nonlinear utility function.

	bet-on-Y	*bet-on-K*	*bet-on-U*	*bet-on-B*
b_{20}	$u(2)\frac{100}{900}+u(1)\frac{400}{900}$	$u(2)\frac{200}{900}+u(1)\frac{500}{900}$	$u(2)\frac{200}{900}+u(1)\frac{500}{900}$	$u(2)\frac{400}{900}+u(1)\frac{400}{900}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_{15}	$u(2)\frac{150}{900}+u(1)\frac{450}{900}$	$u(2)\frac{200}{900}+u(1)\frac{500}{900}$	$u(2)\frac{225}{900}+u(1)\frac{450}{900}$	$u(2)\frac{300}{900}+u(1)\frac{450}{900}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_{10}	$u(2)\frac{200}{900}+u(1)\frac{500}{900}$	$u(2)\frac{200}{900}+u(1)\frac{500}{900}$	$u(2)\frac{200}{900}+u(1)\frac{500}{900}$	$u(2)\frac{200}{900}+u(1)\frac{500}{900}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_5	$u(2)\frac{250}{900}+u(1)\frac{550}{900}$	$u(2)\frac{200}{900}+u(1)\frac{500}{900}$	$u(2)\frac{125}{900}+u(1)\frac{650}{900}$	$u(2)\frac{100}{900}+u(1)\frac{550}{900}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_0	$u(2)\frac{300}{900}+u(1)\frac{600}{900}$	$u(2)\frac{200}{900}+u(1)\frac{500}{900}$	$u(2)\frac{0}{900}+u(1)\frac{900}{900}$	$u(2)\frac{0}{900}+u(1)\frac{600}{900}$

[10]In the risk neutral case, the pure Nash equilibria strategies used by player D correspond to his pure maxmin strategies. This is due to the game being zero-sum. The additional pure Nash equilibria strategies that appear in the benevolent game correspond to his maxmax strategies. This makes intuitive sense: when playing against an adversial player F in the malevolent game, the cautious maxmin strategies are reasonable. In the benevolent game, where F's payoff is maximized by maximizing player D's payoff, using maxmax is.

The best response structure of a risk averse or risk loving player D is identical to the case of risk neutrality: Given b_{10}, all strategies of D are a best response. In the case of fewer black balls, the only best response is *bet-on-Y*. For more black balls, the only best response is to play *bet-on-B*. This implies that the Nash equilibria are not changed by the risk attitude of player D.

However, this is not the case for player F. A player F with a nonlinear utility function has an incentive to deviate from the uniform distribution of balls given that player D chooses *bet-on-U*. The deviation depends on the curvature of his utility function and the treatment. In the benevolent treatment when player F's utility function is (strongly) concave, his best response to *bet-on-U* is to fill the Ellsberg urn with no black balls and 20 yellow balls. However, when player F's utility function is (strongly) convex, the distribution containing 15 black and 5 yellow balls is his best response to *bet-on-U*. In the malevolent treatment, the deviations are the opposite of that. In either case, NE_2 does not exist when player F is not risk neutral. The strategy combination where player D decides to *bet-on-K* and player F chooses the uniform distribution is the *only* pure Nash equilibrium in the malevolent game, NE_1.

Table 5: Pure Nash equilibria - malevolent treatment.

		Player D		
		RA	RN	RL
Player F	RA	NE_1	NE_1	NE_1
	RN	$NE_1; NE_2$	$NE_1; NE_2$	$NE_1; NE_2$
	RL	NE_1	NE_1	NE_1

Summing up, when the Ellsberg task is conceived as an interactive situation, the game-theoretic predictions are different for the malevolent and the benevolent variant of the 3-color Ellsberg game. In the malevolent treatment (see Table 5),

Table 6: Pure Nash equilibria - benevolent treatment.

		Player D		
		RA	*RN*	*RL*
Player F	*RA*	$NE_1; NE_3; NE_4$	$NE_1; NE_3; NE_4$	$NE_1; NE_3; NE_4$
	RN	$NE_1; NE_2; NE_3; NE_4$	$NE_1; NE_2; NE_3; NE_4$	$NE_1; NE_2; NE_3; NE_4$
	RL	$NE_1; NE_3; NE_4$	$NE_1; NE_3; NE_4$	$NE_1; NE_3; NE_4$

the concept of pure Nash equilibrium predicts *bet-on-K* and, if the urn filler is risk neutral, *bet-on-U*. However, *bet-on-U* is not predicted when the urn filler is risk averse or risk loving. In the benevolent treatment (see Table 6), in addition to the Nash equilibria found in the malevolent case, there exist pure coordination Nash equilibria. These are payoff and risk dominant, and correspond to *bet-on-Y* and *bet-on-B*.

Hypothesis 1. *Pure Nash equilibria predictions:*

a) In the malevolent treatment, subjects in the role of player D will bet-on-K or bet-on-U.

b) In the benevolent treatment, subjects in the role of player D will bet-on-K, bet-on-Y, bet-on-B or bet-on-U.

c) bet-on-Y and bet-on-B will be chosen more often in the benevolent treatment than the malevolent treatment.

The idea to analyse the Ellsberg experiments as a game dates back to Schneeweiss (1968, 1973). In those early papers, Schneeweiss already suggested that the Ellsberg tasks can be analysed as a zero-sum game and calculated the maxmin solutions. For the case of risk neutral subjects, the maxmin solutions are identical to the pure Nash equilibrium strategies. However, the solution concepts differ for the case of risk aversion and risk loving. Using maxmin, it is the Ellsberg decision maker's risk attitude that affects which strategy is chosen. Under the Nash equilibrium, it turns out that, instead, it is the urn filler's risk attitude that determines the decision maker's equilibrium strategies. For the case of non-risk-neutrality, the Nash equilibrium is a better solution concept than maxmin, since the game is no longer zero-sum in utilities.[11] Different to our paper, Schneeweiss did not discuss the impact of Ellsberg's experimental conditions on actual experiments.

Kühberger and Perner (2003) conducted four non-incentiviced hypothetical 2-color Ellsberg urn experiments. In hypothetical experiments, it is possible to adhere to Ellsberg's conditions (since the experimenter can ask the subjects to imagine a situation where they are fulfilled). However, in their experiment, Kuehberger and Perner asked their subjects to imagine that they are playing against an opponent or with a partner, breaking Ellsberg's conditions. They apparently were unaware of Schneeweiss' analysis, but discussed their experiments in terms of using a maxmin or a maxmax decision rule.[12] In their experiments one and two, Kuehberger and Perner found a significant difference between the amount of subjects betting on the risky urn in the partner and opponent condition, but this result changes in experiment four, where the experimental procedures were somewhat changed. Subjects' behavior differed more strongly if subjects were asked to imagine that the partner, in addition to the imagined monetary incentives, was their best friend.

3 Uncertainty Averse Preferences

Above, we derived the game-theoretic prediction for the Ellsberg 3-color experiment. In this section, the decision-theoretic predictions are investigated.

There exist well established theories of decision-making under ambiguity, i.e., uncertainty with unknown probabilities. Typically, the theories are formulated in the setup of Anscombe and Aumann (1963). Let S be a finite set of states

[11] This drawback of using maxmin is already mentioned by Schneeweiss (1973).

[12] For an analysis of the 2-color urn experiment in terms of the Nash equilibria, see Appendix 7 & 8.

and let $\Delta(X)$ be the set of all (simple) lotteries over a set of consequences X.[13] Objects of choice are acts $f : S \to \Delta(X)$; i.e., mappings from states to lotteries. Bets are binary acts. A decision maker is characterized by a preference relation $\succcurlyeq$ on $\mathcal{F}$, the set of all acts. Among lotteries, but not necessarily among acts, preferences obey the axioms of expected utility theory. The decision-theoretic approach characterizes a functional $V(\cdot)$, representing $\succcurlyeq$, which specifies how some form of subjective beliefs on S and a utility function $u : X \to \mathbb{R}$ are aggregated when evaluating acts.

The canonical approach is the Maxmin Expected Utility (MEU) model of Gilboa and Schmeidler (1989). In this theory, beliefs are represented by a non-empty, closed, and convex set $\mathcal{P}$ of probability distributions on S with respect to which acts are evaluated via the following functional:[14]

$$V^{min}(f) = \min_{p \in \mathcal{P}} \int_S u(f(s))dp(\{s\}) \qquad \forall f \in \mathcal{F}. \tag{2}$$

Following Ghirardato, Maccheroni, and Marinacci (2004), the set of priors $\mathcal{P}$ and the *min*-operator are reflecting the decision maker's perception of ambiguity and reaction to it, respectively. That is, $\mathcal{P}$ represents the probabilistic scenarios that the decision maker views as possible. The *min*-operator encapsulates aversion towards ambiguity. When the *min*-operator is replaced by its dual form, the *max*-operator, the decision maker exhibits proneness towards ambiguity. If the subjective set of priors is a singleton, the operators *min* and *max* are irrelevant and preferences are of the expected utility form. For a given set of priors $\mathcal{P}$ and an operator $k \in \{min, max\}$, V^k denotes the k-functional representing $\succcurlyeq$ on $\mathcal{F}$.

Let us be more specific about how the MEU model rationalizes choices in the 3-color experiment. The states of nature correspond to the outcome of a random

[13] The Anscombe-Aumann framework is a special case of the "purely subjective" setup of Savage (1954) in which the set of consequences is an arbitrary set.

[14] The MEU model has been originally axiomatized by Gilboa and Schmeidler (1989) in the setup of Anscombe and Aumann (1963). Alternative axiomatizations, in the subjective setup of Savage (1954), were provided by Casadesus-Masanell, Klibanoff, and Ozdenoren (2000a,b); Ghirardato, Maccheroni, Marinacci, and Siniscalchi (2003); and Alon and Schmeidler (2014).

draw, the color of the ball. That is, $S = \{R, B, Y\}$. The bets of the choice problems I and II (see Table 1) are depicted as mappings from S to degenerate lotteries yielding \$1 or \$0, respectively. We assume that each decision maker incorporates the probabilistic information available in the 3-color experiment. That is, $p(R) = \frac{1}{3}$ for any $p \in \mathcal{P}$. For a given set of priors $\mathcal{P}$, denote by $\underline{p}(E) = \min_{p \in \mathcal{P}} p(E)$ the lowest probability of an event $E \in \{B, Y\}$. Similarly, $\overline{p}(E) = \max_{p \in \mathcal{P}} p(E)$. Clearly, if $\underline{p}(E) = \overline{p}(E)$ for an event $E \in \{B, Y\}$, the set of priors $\mathcal{P}$ is a singleton with a probability distribution denoted by p^*.

Any pair $\mathcal{P}$ and V^k characterizes a subject's betting behavior in the Ellsberg-task. Let the utility function be normalized to $u(\$0) = 0$ and $u(\$1) = 1$. Then, a decision maker with $\mathcal{P}$ and V^k chooses[15]

$$\begin{array}{llcll}
(i) & \textit{bet-on-K} & \Longleftrightarrow & \frac{1}{3} > \underline{p}(B) \quad \text{and} & \frac{1}{3} > \underline{p}(Y), \\
(ii) & \textit{bet-on-B} & \Longleftrightarrow & \frac{1}{3} < p^*(B) \quad \text{and} & \frac{1}{3} > p^*(Y), \\
(iii) & \textit{bet-on-Y} & \Longleftrightarrow & \frac{1}{3} > p^*(B) \quad \text{and} & \frac{1}{3} < p^*(Y), \\
(iv) & \textit{bet-on-U} & \Longleftrightarrow & \frac{1}{3} < \overline{p}(B) \quad \text{and} & \frac{1}{3} < \overline{p}(Y).
\end{array}$$

The crucial feature of the MEU model is that both choice problems I and II are evaluated *separately*. Consider a decision maker who has a non-singleton $\mathcal{P}$ and chooses *bet-on-K*. Although there is one urn, the decision maker believes that two distinct probability distributions are possible, one in each problem:

$$p_1 = \left(\frac{1}{3}, \underline{p}(B), \frac{2}{3} - \underline{p}(B)\right) \text{ and } p_2 = \left(\frac{1}{3}, \frac{2}{3} - \underline{p}(Y), \underline{p}(Y)\right).$$

In problem I, the expected utility of f_2 is the smallest one under p_1 and thus p_1 is the probability distribution rationalizing the decision maker's preference for f_1. Put differently, he believes that p_1 is the true distribution from which a random draw will be taken to determine his payoff. For the same reason, the decision maker considers p_2 as the true distribution when comparing acts f_3 and f_4 in problem II.

This is exactly the manner in which the "malevolent Nature" enters into the model via the *min*-operator: Nature is malevolent by always, "magically", drawing

[15] Cases (ii) and (iii) follow from the fact that $\underline{p}(B) = \frac{2}{3} - \overline{p}(Y)$ whenever $p(R) = \frac{1}{3}$ $\forall p \in \mathcal{P}$.

a ball from such a distribution that is the worst possible for this very part of the problem at which the decision maker is currently looking. Under the max-operator, the logic works in the opposite direction and Nature is perceived as benevolent. Given the max-operator and set $\mathcal{P}$, the decision maker behaves ambiguity loving and chooses f_2 in problem I and f_3 in II. As mentioned before, if the subjective set $\mathcal{P}$ is a singleton, the min- and max-operator play no role and the decision maker behaves ambiguity neutral. In particular, for a nonuniform prior, the decision maker chooses either f_1 and f_3 or f_2 and f_4 (i.e. *bet-on-*Y or *bet-on-*B).

In all treatments, subjects receive the same information about the Ellsberg urn. Thus, it is assumed that a subject who would perceive ambiguity (i.e., would have a set of priors $\mathcal{P}$ with $\underline{p}(B) \neq \overline{p}(B)$) in treatment M, would also perceive ambiguity in treatment B. Yet, the subject may have different sets of priors in both treatments as long as they are non-singletons. Analogously, a subject who would perceive no ambiguity (i.e., would have a single prior with $p^*(B) \in [0, \frac{2}{3}]$) in one treatment would not perceive ambiguity in the second treatment. Again, such a subject may have different singletons in each treatment. Since the k-operator is a component of the subjects' preferences, we further assume that their attitude towards the perceived ambiguity remains the same among treatments. Under these assumptions, the decision-theoretic approach predicts no treatment effect. That is, one expects to observe the same amount of subjects who *bet-on-K*, *bet-on-U*, and *bet-on-C* in all treatments.

Hypothesis 2. *Uncertainty averse preference predictions:*

a) Consider a subject with preferences represented by V^k with respect to a set $\mathcal{P}$. Depending on set $\mathcal{P}$ and the k-operator, the subject in the role of player D will choose either bet-on-K, bet-on-U, or bet-on-C.

b) The frequency of subjects who bet-on-K, bet-on-U and bet-on-C should not change between treatments.

As mentioned previously, the MEU model is said to admit a "game against Nature" interpretation. That is, an ambiguity averse decision maker is interpreted as playing a hypothetical zero-sum game against a "malevolent Nature".[16] In this case, an act f and a probability distribution p become, respectively, the strategies of the decision maker and of Nature (see Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio, 2011). For instance, if the decision maker has MEU preferences and contemplates choosing an act f, Nature responds by choosing a probability from a set $\mathcal{P}$ so as to minimize the decision maker's expected utility given his choice of act f (see Equation 2).

Bearing in mind this interpretation, one could assume that for a given set of priors $\mathcal{P}$, the k-operator is correlated with the treatment, i.e., with the motives of the urn filler (player F). Since the motives are explicitly stated in each treatment, it could be that subjects who perceive ambiguity (i.e., with $\underline{p}(B) \neq \overline{p}(B)$) evaluate acts according to V^{min} in the malevolent treatment and according to V^{max} in the benevolent treatment. In this case, no *bet-on-U* should be observed in the malevolent treatment and no *bet-on-K* should be observed in the benevolent treatment. For subjects whose sets of priors are singletons, the *min*- and *max*- operators are irrelevant and this distinction does not matter. Therefore, the same amount of subjects who *bet-on-C* should be observed in both treatments.

Although both the "game against Nature" interpretation and the 3-color Ellsberg game feature the aspect of playing against an opponent, there is a strong conceptual difference. According to the "game against Nature" interpretation of UAP, when taken literally, the Ellsberg task actually consists of two, *separate*, zero-sum games of the subject versus Nature. One game for the first decision and

[16] Based on the "game against Nature" interpretation, Ozdenoren and Peck (2008) use the notion of subgame perfection to address the issue of dynamic consistency among ambiguity averse decision makers in a dynamic version of the 3-color Ellsberg task. Similar to the UAP framework, they keep the two questions separated and only consider responses to the second question as strategies. A real experiment on a dynamic version of the 3-color experiment was conducted by Dominiak, Duersch, and Lefort (2012), finding that subjects behave in line with consequantialism rather than dynamic consistency. One of these fundamental properties of dynamic behavior has to be relaxed when modeling ambiguity aversion in dynamic setups (see Ghirardato, 2002; Siniscalchi, 2011).

a second game for the second decision. In contrast, the 3-color Ellsberg game is a *single game* featuring all actions by the decision maker as his strategy. As such, there is no separation of the two questions asked during the Ellsberg task. Therefore, malevolence is captured very differently in the game-theoretic setup. It is not modeled via the functional representing the decision maker's preferences, as in the UAP model, but, via the utility function of the urn filler. The structural differences between the two approaches are explored in more detail in Section 4.

There is a broader class of ambiguity models admitting the "separation-principle", the "game against Nature" interpretation, and thus supporting Hypothesis 2. It is the family of uncertainty averse preferences (UAP) introduced by Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio (2011).[17] Let $\Delta(S)$ be the set of all probability distributions over the state space S. The UAP are represented by

$$V^{UA}(f) = \min_{p \in \Delta(S)} G\Big(\int u(f) dp, p \Big) \qquad \forall f \in \mathcal{F}, \tag{3}$$

where $G : \mathbb{R} \times \Delta(S) \to (-\infty, \infty]$ is a quasiconvex function, increasing in the first variable, and satisfies $\inf_{p \in \Delta(S)} G(t, p) = t$ for all t. Function G is interpreted as an ambiguity aversion index. Let $c : \Delta(S) \to [0, \infty]$ be a convex (cost) function on $\Delta(S)$ satisfying $\inf_{p \in \Delta(S)} c(p) = 0$. When $G(t, p) = t + c(p)$, V^{UA} represents the variational preferences of Maccheroni, Marinacci, and Rustichini (2006). For a given set of priors $\mathcal{P}$, V^{UA} with $G(t, p) = t + c(p)$ represents MEU preferences whenever $c(p) = 0$ for any $p \in \mathcal{P}$ and otherwise $c(p) = \infty$. The family of UAP unifies all the generalizations of the MEU model that maintain the key axiom of Schmeidler (1989), called Uncertainty Aversion.[18]

[17] Another model supporting Hypothesis 2 is the smooth ambiguity model of Klibanoff, Marinacci, and Mukerji (2005). The model is built on the "two-stage approach" dating back to Segal (1987). A decision maker has a subjective second order prior: that is, a prior over the possible compositions of the Ellsberg urn. A concave (resp. convex) second order "utility function" (the so-called ϕ-function) corresponds to ambiguity averse (resp. loving) behavior. The prediction of player D's betting behavior is driven by the curvature of his ϕ-function.

[18] The axiom is the behavioral counterpart of an averse attitude towards ambiguity. Roughly, it

4 Comparing the Theories

What is the structural difference between the decision-theoretic models and the game-theoretic one? Each model makes different assumptions about the characteristics of the situation that the subjects are supposed to take into account when making their decision. As a general rule, the UAP models restrict the subjects to consider the outcome of a random draw from the urn as the only payoff-relevant source of uncertainty. The decision is derived from the decision maker's tastes over outcomes in addition to some form of subjective beliefs. The form of the beliefs might differ depending on the model, but is supposed to encompass the ambiguity perceived by the decision maker (Ghirardato, Maccheroni, and Marinacci, 2004). Not taken into account are the motivations of the person filling the urn. The situation is not modeled as a strategic (interactive) one.

In the game-theoretic model, it is assumed that both players are expected utility maximizers. That is, the model strongly restricts the form of players' preferences. However, it additionally allows subjects to take into account a broader amount of information about the situation they face. Instead of only looking at possible compositions of the Ellsberg urn, subjects now consider as well the strategic implications of how the urn was set up. As is shown in Section 2, even when using the most standard way of modeling preferences, taking into account the strategic implications of the decision situation leads to a prediction of *bet-on-K* or *bet-on-U* in the 3-color Ellsberg experiment.[19] Note that, since the Nash model does not possess free parameters, we arrive at this prediction via a model that is more parsimonious compared to the ambiguity models discussed in Section 3.

---

requires that a decision maker always displays a weak preference for mixtures (randomizations) between two acts among which the decision maker is indifferent. Recently, Dominiak and Schnedler (2011) provided an experimental evidence showing that ambiguity averse subjects rather behave neutral towards mixtures, or even dislike randomization between two ambiguous acts.

[19] Several papers generalize the Nash equilibrium notion to ambiguity-sensitive preferences: Dow and Werlang (1994), Lo (1996), Eichberger and Kelsey (2000), Marinacci (2000), Bade (2011) and Lehrer (2012).

By far, the majority of Ellsberg urn experiments feature payoffs that correspond to the malevolent game. That is, each amount of money paid to subjects is paid out of the pockets of the person or institution which is also responsible for setting up the composition of the urn. Subjects who take into account the monetary motivations of the experimenters can legitimately consider the situation to be adversary.

Table 7: Frequency predictions of behavior for players D.

UAP Models			Game Theory		
Malevolent Treatment		Benevolent Treatment	Malevolent Treatment		Benevolent Treatment
$\#\underbrace{\text{bet-on-K}}_{\text{“}averse\text{”}}$	$\geq$	$\#\underbrace{\text{bet-on-K}}_{\text{“}averse\text{”}}$	$\#\underbrace{\text{bet-on-K}}_{\mathbf{NE_1}}$	$\geq$	$\#\underbrace{\text{bet-on-K}}_{\mathbf{NE_1}}$
$\#\underbrace{\text{bet-on-U}}_{\text{“}loving\text{”}}$	$\leq$	$\#\underbrace{\text{bet-on-U}}_{\text{“}loving\text{”}}$	$\#\underbrace{\text{bet-on-U}}_{NE_2}$		$\#\underbrace{\text{bet-on-U}}_{NE_2}$
$\#\underbrace{\text{bet-on-C}}_{\text{“}neutral\text{”}}$	$=$	$\#\underbrace{\text{bet-on-C}}_{\text{“}neutral\text{”}}$	#bet-on-C	$\leq$	$\#\underbrace{\text{bet-on-C}}_{\mathbf{NE_3\&NE_4}}$

UAP: The number of subjects who bet-on-C is equal across treatments. The fraction of subjects who bet-on-K (bet-on-U) is larger (smaller) or equal in malevolent than in benevolent. GT: The number of subjects who bet-on-C is higher in Benevolent. (Note: Nash equilibrium NE_2 only exists if player F is risk-neutral.)

The game-theoretical interpretation of the Ellsberg task delivers a prediction of treatment differences in our experiment: more subjects D are predicted to choose *bet-on-C* in the benevolent game compared to the malevolent game (see Table 7). The ambiguity models can predict all choices, but only when the perceived ambiguity and attitudes change across treatments. Changes in priors change whether subjects *bet-on-C* or not, given that the set of priors changes from a singleton to a non-singleton. However, the prediction stays unchanged when priors differ but stay a singleton or stay a non-singleton. Changing the *min*-operator to into a *max*-operator, changes whether subjects *bet-on-K* or *bet-on-U*. If the percieved ambiguity and the operator do not change the ambiguity models would not predict a difference in the number of subjects who *bet-on-C*, *bet-on-K* and *bet-on-U* between the malevolent and benevolent treatments.[20]

We conduct an experiment to empirically test the different predictions made with actual behavior by subjects in the malevolent and benevolent case.

5 Experiment

5.1 Experimental Design

The experimental design includes three treatments - **(B)**enevolent, **(M)**alevolent and **(S)**tandard - which all employ the 3-color Ellsberg urn to elicit subjects' behavior in the face of an ambiguous situation.

In treatments **B** and **M**, the Ellsberg decisions are made by subjects of type D, while the urns are filled by subjects of type F. Both treatments are identical up to the monetary incentives of subjects F. At the start of the experiment, each subject is randomly assigned role D or role F, and Ds and Fs are anonymously matched into pairs. Each subject F is separately tasked with determining the distribution of a 3-color Ellsberg urn. To this end, all Fs receive a nontransparent urn already filled with 10 red balls. Additionally, they receive 20 black and 20 yellow balls. Out of these, they are asked to choose 20 balls, in any combination, to add to the 10 red balls. The thus constructed 3-color urns are taken by the experimenter and placed on a table in sight of all subjects.[21]

Subjects of type D are asked to make the usual 3-color Ellsberg choices. That

[20]There is no within-model way of obtaining the piors or the operator from the description of the situation faced by the decision maker. Note that a model which can predict each possible outcome (given some parameters), and does not tie down parameters, essentially has no predictive power. Regarding our main question, whether the number of subjects who *bet-on-C* changes across treatments, we assume than an equal number of subjects exhibits singleton priors and non-singleton priors in each treatment.

[21]The remaining balls are used by the experimenter to fact check that all urns contain the correct number of balls.

is, each D is asked to make a decision on the following two choices, each refering to a separate draw (with replacement), conducted later, out of the Ellsberg urn constructed by the matched F:

- **First choice:** Do you prefer to
  - receive a payment of 4 EUR if the drawn ball is red, or
  - receive a payment of 4 EUR if the drawn ball is black?
- **Second choice:** Do you prefer to
  - receive a payment of 4 EUR if the drawn ball is red or yellow, or
  - receive a payment of 4 EUR if the drawn ball is black or yellow?

Note that we use the standard framing of the Ellsberg task, where two questions are asked separately, instead of a possible different framing where strategies are elicited via a single question (with four alternatives). As such, our design is conservative when testing whether subjects view the two questions as one integrated task. Subjects D are also asked to indicate their confidence in their own decisions on a five point Likert scale.[22]

In treatment **B**, the incentives of D and F are perfectly aligned: whatever D wins in the Ellsberg task will be paid to F as well. In treatment **M**, this is reversed. F is paid a lump sum of 8 EUR, and all earnings by D are deducted from this.[23] Thus, treatments **B** and **M** implement the benevolent and malevolent Ellsberg game discussed in Section 2.

Treatment **S** corresponds to the standard Ellsberg procedures. All subjects are of type D and the urn was filled by the experimenter. The instructions gave no specific procedure indicating how the urn was filled.

---

[22]See http://www.adominiak.com/InstructionsDominiakDuersch.zip for all instruction material.

[23]If both D and F choose randomly, the expected earnings from the task in both treatments are 4 EUR for all subjects.

All subjects received a show up fee of 4 EUR and answered a questionnaire at the end of the experiment. As part of the unpaid questionnaire, subjects were asked for a point estimate of the number of black and yellow balls in the urn if they were of type D. Both D and F had to answer how they would have, hypothetically, chosen if placed in the other role. Finally, the questionnaire included a 10 item optimism/pessimism scale.[24]

5.2 Procedures

The experiment consisted of eleven sessions and was conducted at the laboratory of University of Heidelberg from June 2010 to March 2011 and in May 2013. All recruitment was done via ORSEE (Greiner (2004)). In total, 169 subjects took part in our experiment, 88 male and 81 female.[25]

While entering the lab, subjects randomly picked a table tennis ball labeled with their seat number and a letter. In treatments **B** and **M**, the letter was used to match subjects D and F, who were seated at opposite walls of the lab. No communication was allowed, and we did not observe any instances of attempted communication by subjects. The matching rule was explained in neutral wording by the experimenter before the instructions were handed out. Subjects were told that everyone received the same instructions. Questions were answered in private by the experimenter. After Fs had filled their urns and Ds marked their decisions on the decision sheet, the sheets and urns were collected by the experimenter.[26]

The draws from each urn were done by the experimenter in public in the middle of the room. Afterwards, the experimenter calculated the earnings of each subject in private, while subjects answered the questionnaire. Finally, all subjects were

[24]German version of the Life-Orientation-Test, LOT-R (Glaesmer, Hoyer, Klotsche, and Herzberg, 2008).

[25]One subject F has been excluded from the dataset, since he had participated in two sessions.

[26]Urns were collected first to ensure that Ds could make their decision after the composition of the urn was determined. Both urns and decision sheets were labeled with the same letter as the respective table tennis ball.

called according to their seat number and paid their earnings in private and in cash. Subjects earned on average 7.50 EUR for roughly 45 minutes in the lab (7.56 EUR in **B**, 6.60 EUR in **S**, and, by construction, 8.00 EUR in **M**).

In treatment **S**, no subjects of type F took part. All participants were of type D. A single urn was used and placed in sight of the subjects for the entire experiment.

5.3 Results

We start by describing the behavior of the 104 subjects D, our main group of interest. Averaged over all treatments, 8.7% of subjects *bet-on-U*, 44.2% *bet-on-K* and 47.1% *bet-on-C*.[27] However, as depicted in Figure 1, this varies across treatments: while in treatment **B** 22.6% of subjects D *bet-on-K*, this jumps to 48.5% in treatment **M** and 57.5% in treatment **S**. Using the Fisher exact test, the differences between **B** and **M** (two-sided, p=.039, obs=64) and between **B** and **S** (two-sided, p=.004, obs=71) are significant. However, **M** and **S** are not significantly different (two-sided, p=.486, obs=74). This result is mirrored when looking at subjects who *bet-on-C*. Treatments **B** and **M** are significantly different (two-sided, p=.047, obs=64) and so are **B** and **S** (two-sided, p=.001, obs=71), while **M** and **S** are not (two-sided, p=.225, obs=74). The share of subjects who *bet-on-U* is never significantly different between any two treatments.

Result 1a: More subjects choose *bet-on-K* and fewer subjects *bet-on-C* in the malevolent treatment compared to the benevolent treatment.

Result 1b: In all treatments, *bet-on-U* is chosen least often.

Concerning the across treatment predictions, we thus find evidence in favor of the Nash prediction (Hypothesis 1c), while the results are not compatible with the UAP prediction (Hypothesis 2b). Regarding within treatment predictions, the UAP theories can explain every possible betting behavior. Hypothesis 2a is

[27]Out of the subjects who *bet-on-C* , 53.1% favored black and 46.9% favored yellow.

therefore trivially satisfied in all treatments. Pure Nash equilibrium Hypothesis 1b for treatment **B** is similar. In the case of risk neutral players F, every betting behavior can be explained as a pure Nash equilibrium strategy. If players F are not risk neutral, *bet-on-*U is not predicted. In the experiment, there are indeed very few subjects who *bet-on-*U. Hypothesis 1a for treatment **M** is not confirmed, since there exist players D in **M** who *bet-on-*C.

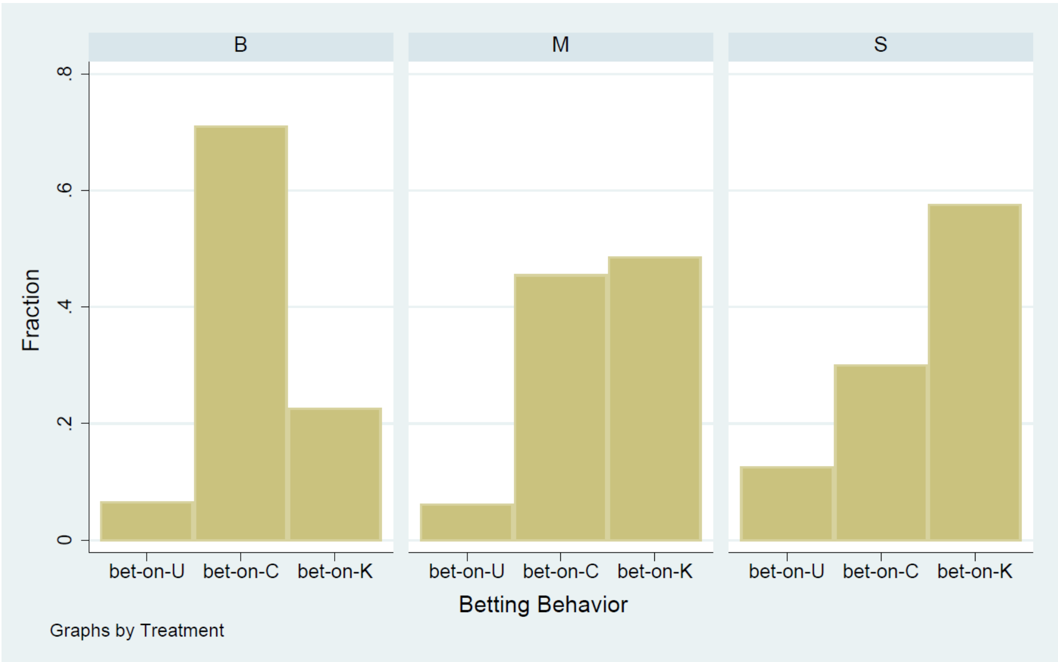

Figure 1: Betting behavior by treatment.

Behavior in treatment **S** suggests that subjects perceive the standard implementation of the 3-color Ellsberg task as a game against the experimenter and view the experimenter as an opponent rather than as an ally. There are no significant differences between **M** and **S**, and, if anything, subjects seem to be even more inclined to *bet-on-*K in **S**. In most Ellsberg experiments, the filling of the urn, and the monetary consequences for subjects and experimenters are set up similar to our treatment **S**. That is, a large part of the experimental literature looks only at one case of implementing the Ellsberg task, namely subjects reacting to a malevolent urn filler. The other case, reacting to a benevolent urn filler, has been ignored so far.[28]

Result 2: Betting behavior in the *standard* treatment is significantly different from the *benevolent* treatment.

Table 8 shows a multinomial probit regression controlling for being female, participation in a formal course in game theory or statistics, optimism, whether the subjects play games of chance; and the subjects' self-reported confidence in their choices. The regression compares subjects who *bet-on-C* and *bet-on-U* to the base outcome, *bet-on-K*. Result 1 is confirmed. Subjects in the benevolent treatment choose significantly more often *bet-on-C* rather than *bet-on-K*, compared to the malevolent treatment. This is not true for the standard treatment, where the difference is insignificant. Gender or playing games of chance do not have a significant impact; neither does being more optimistic on the psychological optimism scale we elicit in the questionnaire nor having participated in a formal course in game theory or statistics. However, subjects who choose *bet-on-K* have a higher average confidence in their choices.[29]

Up till now, we evaluated subjects D, who choose more *bet-on-C* in treatment **B**. That is, they tried to coordinate on one color. Now, we look at subjects F. Do these, as well, try to coordinate? As Figure 2 shows, they do. The number of subjects F playing a pure coordination strategy (no black balls at all, or only black balls, corresponding to their strategy in NE_3 and NE_4) is significantly higher in the benevolent treatment (Fisher exact test, two-sided, p=0.023, obs=65). That is, both players try to achieve the higher coordination payoff in treatment **B**, but shy away from this potentially exploitable strategy when in a zero-sum payoff game in

[28]See Camerer and Weber (1992) and Trautmann and van de Kuilen (2014) for a review of ambiguity experiments.

[29]Overall, our subjects are reasonably confident in their choices. On average, they report a confidence of 2.45 and 2.95 (out of $\{0, 1, 2, 3, 4\}$) for the first and the second task. Only 6, respectively 3, subjects reported having a confidence of zero. All of our results are unchanged if we exclude subjects with zero confidence from the analysis.

Table 8: Multinomial probit on betting behavior.

Multinomial Probit Regression				Number of obs	=	104
				Wald $\chi^2(14)$	=	25.33
Log Likelihood = -82.075347				Prob > χ^2	=	0.0314

Betting Behavior	Coef.	Std. Err.	z	$P >\mid z\mid$	95% Conf. Inter.	
bet-on-U						
Treatment S	0.5473915	0.7228561	0.76	0.449	-0.869380	1.964163
Treatment B	0.5214416	0.7750476	0.67	0.501	-0.997624	2.040507
Optimism	0.0842592	0.0553064	1.52	0.128	-0.024139	0.192658
AVGconfidence	-0.7010100	0.3283717	-2.13	0.033	-1.344607	-0.057414
Female	0.3138005	0.5968215	0.53	0.599	-0.855948	1.483549
Formal Course	-0.6276003	0.6304215	-1.00	0.319	-1.863204	0.608003
Games of Chance	-0.1157415	0.7190218	-0.16	0.872	-1.524998	1.293515
Cons	0.2782610	0.9929528	0.28	0.779	-1.667891	2.224413
bet-on-C						
Treatment S	-0.3940401	0.4846323	-0.81	0.416	-1.343902	0.555822
Treatment B	1.0592460	0.4821105	2.20	0.028	0.114327	2.004166
Optimism	0.0055134	0.0347493	0.16	0.874	-0.062594	0.073621
AVGconfidence	-0.4681461	0.2309238	-2.03	0.043	-0.920748	-0.015544
Female	-0.1847030	0.4052473	-0.46	0.649	-0.978973	0.609567
Formal Course	0.4485750	0.4078894	1.10	0.271	-0.350874	1.248024
Games of Chance	0.4201125	0.4603592	0.91	0.361	-0.482175	1.322400
Cons	0.9027201	0.7178993	1.26	0.209	-0.504337	2.309777
bet-on-K	(base outcome)					

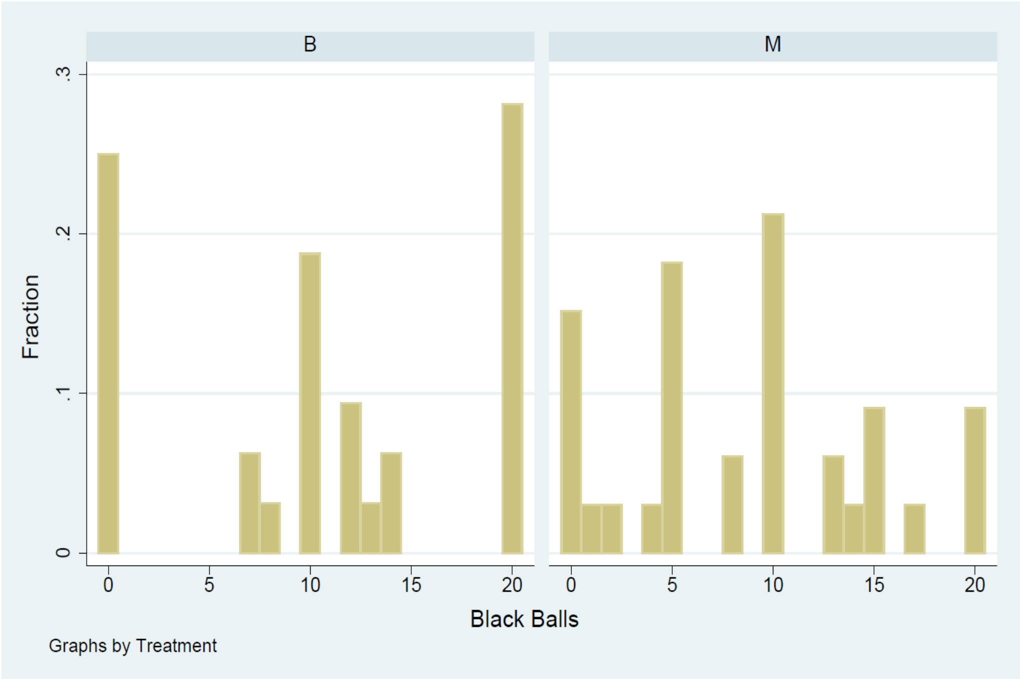

Figure 2: Histogram of black balls used by subjects F.

treatment **M**. This is in line with the pure Nash equilibrium prediction.

Result 3: Subjects F attempt to coordinate more in treatment **B** than in treatment **M**.

A last question concerns the success of coordination. Did subjects manage to coordinate on one color? The results show that they did not. While both groups of subjects tried to coordinate, they failed to pick one color on which to coordinate. 62.5% of subjects D who attempted to coordinate did so on the strategy favoring black; but out of the coordinating F, only 52.9% used black balls. Therefore subjects reap only a very small advantage from their coordination attempts.[30]

Result 4: Coordination on one color is attempted, but fails.

[30]Due to unlucky matching and draws, subjects in treatment **B** achieved an average payoff even slightly less than the expected one for completely random choice on both sides.

6 Concluding Remarks

Experimental implementations of Ellsberg's thought-experiments have traditionally been used as evidence against SEU and in favor of models of ambiguity-sensitive preferences. We design an experiment showing that, rather than displaying ambiguity-sensitive preferences in individual choice problems, subjects in Ellsberg experiments are better described as playing a game against the person who determined the distribution of balls in the urn. In two treatments, which keep all information about the Ellsberg urn unchanged - while varying the incentives of the urn filler - we find a treatment difference as predicted by the pure Nash equilibrium, but not predicted by the MEU model or its extensions.

While our work is concerned with the conclusions typically drawn from Ellsberg experiments, Al-Najjar and Weinstein (2009a,b) challenged the existing ambiguity models by questioning the rationality of ambiguity averse behavior.[31] They also criticized the significance of the models as a descriptive approach. As a main argument, Al-Najjar and Weinstein argued that the Ellsberg-type behavior can be consistently explained by heuristics learned by subjects in the real-word context but misapplied in the laboratory environment: "*the basic experimental findings supporting Ellsberg choices cannot distinguish between two competing explanations: 1. Ambiguity models which explain these choices by appealing to taste (ambiguity aversion). 2. A model where subjects misapply heuristics that serve them well in real-world situations.*" (Al-Najjar and Weinstein, 2009a, p. 277)

This paper demonstrates that there is no need to appeal to heuristics. Solving the Ellsberg task as one game, using standard expected utility preferences, delivers predictions that explain the observed behavior. Moreover, creating an actual game with two subjects as players allows for an experimental test of the two com-

[31] As an argument against the rationality of the ambiguity models, Al-Najjar and Weinstein (2009a) scrutinize some behavioral anomalies displayed by ambiguity averse subjects facing a dynamic decision problem; e.g., the sunk cost fallacy or aversion towards new information. For counterarguments and further discussion on the topic, see Mukerji (2009); Siniscalchi (2009).

peting theories: game theory versus ambiguity models. Using a non-hypothetical experiment, we are not restricted to argue our point from the practicality and reasonableness of the theories alone. The experimental results show different behavior under different incentives of the opponent, contrary to the predictions of the UAP models. On the other hand, the viewing the task as a game explains the treatment differences. This provides evidence that subjects interpret the Ellsberg task as a game against the filler of the urn and not as isolated individual decision tasks, as implied by the conditions placed by Ellsberg on his thought-experiments.

One caveat is in order. In the *malevolent* treatment, 45.5% of the subjects *bet-on-C*, even though this is not a pure Nash equilibrium. While *bet-on-C* can be part of mixed Nash equilibria, it seems implausible to us that such a high proportion of subjects would use mixed Nash equilibria. It is likely that not all subjects analyse the situation as a game and integrate their choices into one decision. A part of the subjects may indeed ignore the interaction with the urn filler, as assumed by Ellsberg and the ambiguity literature. However, note that this cannot be interpreted as evidence of ambiguity aversion. The observations which cannot be explained by pure Nash equilibria strategies are those who *bet-on-C*. That is, if these unexplained subjects treat the situation as a one-person decision problem, then they act ambiguity neutral.

What is the consequence of our findings? Up until now, the Ellsberg experiment results have typically been interpreted as strong evidence against the standard SEU theory, and were, therefore, the empirical underpinning of a large amount of literature on preference-based theories which were developed to resolve the Ellsberg paradox. This study has demonstrated that Ellsberg's results are not paradoxical when the task is viewed as a game by the subjects. As such, SEU cannot be rejected unless all the Ellsberg conditions are satisfied. There is a new need to develop experimental tests of SEU that take Ellsberg's conditions into account, if possible at all.

7 Appendix: The 2-Color Experiment

In the 2-color version of Ellsberg's thought-experiment, 2 urns are used. A known urn (K) containing 50 red and 50 black balls, and an unknown urn (U), containing 100 balls in total, with an unknown number of red or black balls. Subjects are asked, for each color, whether they want to have a draw (to win a prize) from the known or the unknown urn. Call r_K a bet on drawing from the known urn, paying \$1 if the draw is red and \$0 otherwise. Bets r_U, b_K, b_U are defined equivalently. Then, the strategy set of player D is $S_D = \{(r_K, b_K); (r_K, b_U); (r_U, b_K); (r_U, b_U)\}$. The strategy set of player F is $S_F = \{b_0, \ldots, b_{100}\}$. Suppose first that both players are risk neutral. Table 9 represents the players' expected payoffs in the malevolent 2-color Ellsberg game. The two pure Nash equilibria are: $NE_1 = \{b_{50}; (r_K, b_K)\}$ and $NE_2 = \{b_{50}; (r_U, b_U)\}$. That is, choosing *bet-on-K* or *bet-on-U* by player D is a pure Nash equilibrium. In either case, the run filler uses exactly 50 balls of each color.

Table 9: Expected payoffs in the malevolent 2-color Ellsberg game.

	$\boldsymbol{(r_K, b_K)}$	(r_K, b_U)	(r_U, b_K)	$\boldsymbol{(r_U, b_U)}$
b_{100}	$-1, 1$	$-\frac{3}{2}, \frac{3}{2}$	$-\frac{1}{2}, \frac{1}{2}$	$-1, 1$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\boldsymbol{b_{50}}$	$\boldsymbol{-1, 1}$	$-1, 1$	$-1, 1$	$\boldsymbol{-1, 1}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_0	$-1, 1$	$-\frac{1}{2}, \frac{1}{2}$	$-\frac{3}{2}, \frac{3}{2}$	$-1, 1$

Next, suppose that player D has a nonlinear utility function. The corresponding expected payoffs are summarized in Table 10. If player F is risk neutral, the Nash equilibria remain unchanged. However, if player F is risk averse with a (strongly) concave utility function then the only Nash equilibrium is $NE_1 = \{b_{50}; (r_K, b_K)\}$.

In a benevolent version of the 2-color Ellsberg game (not depicted here), additional pure Nash equilibria arise where the urn filler uses only one color, on which

Table 10: Expected payoffs of D in the 2-color Ellsberg game - nonlinear utility.

	(r_K, b_K)	(r_K, b_U)	(r_U, b_K)	(r_U, b_U)
b_{100}	$u(2)\frac{1}{4}+u(1)\frac{1}{2}$	$u(2)\frac{1}{2}+u(1)\frac{1}{2}$	$u(1)\frac{1}{2}$	$u(1)$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_{50}	$u(2)\frac{1}{4}+u(1)\frac{1}{2}$	$u(2)\frac{1}{4}+u(1)\frac{1}{2}$	$u(2)\frac{1}{4}+u(1)\frac{1}{2}$	$u(2)\frac{1}{4}+u(1)\frac{1}{2}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_0	$u(2)\frac{1}{4}+u(1)\frac{1}{2}$	$u(1)\frac{1}{2}$	$u(2)\frac{1}{2}+u(1)\frac{1}{2}$	$u(1)$

the decision maker bets. The two coordination Nash equilibria of the benevolent game are not affected by the players' risk attitudes.

Summing up, the predictions for the 2-color Ellsberg game mirror the predictions derived in Section 2 for the 3-color Ellsberg game. In the malevolent game, the pure Nash equilibria strategies correspond to betting on the known probability colors or to betting on the unknown probability colors. When the urn filler has a nonlinear utility function, only betting on the known probability colors survives. In the benevolent game, additional coordination equilibria arise.

Sometimes, the 2-color Ellsberg experiment is conducted with only one question. In this shorter version, subjects are asked to bet on one color, not the other. Then betting on the known probability color and a uniform distribution of balls constitute the only pure Nash equilibrium.

8 Appendix: Suspicion

Regarding the 2-color Ellsberg experiment, the following line of reasoning is used against fixing the colors of the bets. The experimenter fills the urn. Subjects are suspicious of the experimenter and therefore do not bet on the (fixed) color in the unknown urn. They reason that the experimenter would not put the color in the urn that subjects can bet on. As a precaution against this suspicion, subjects are not only allowed to choose whether to bet on the risky or the unknown urn, but allowed to choose which color to bet on, as well.

How would the pure Nash equilibria of the 3-color Ellsberg game be impacted if subjects, additionally to choosing the bets, could also choose whether to bet on the unknown probability event suggested by the experimenter - black in problem I and non black in problem II - or to reverse the payoffs on the unknown colors and bet on the unknown color not originally suggested by the experimenter (i.e., yellow in problem I and non yellow in problem II)? Call the former option **K**eep and the latter **S**witch. Consequently, the strategy set of player D enlarges to $S_D = \{(f_1^K, f_3^K); (f_1^K, f_4^K); (f_2^K, f_3^K); (f_2^K, f_4^K); (f_1^S, f_3^S); (f_1^S, f_4^S); (f_2^S, f_3^S); (f_2^S, f_4^S)\}$. For instance, a player choosing (f_2^S, f_3^S) bets on an unknown probability event, but not on black in problem I and non black in problem II, as initially offered. Instead, the player switches and bets on yellow in problem I and non yellow in problem II.

Table 11: Expected payoffs in malevolent 3-color Ellsberg game with switching.

	(f_1^K, f_3^K)	$\boldsymbol{(f_1^K, f_4^K)}$	$\boldsymbol{(f_2^K, f_3^K)}$	(f_2^K, f_4^K)	(f_1^S, f_3^S)	$\boldsymbol{(f_1^S, f_4^S)}$	$\boldsymbol{(f_2^S, f_3^S)}$	(f_2^S, f_4^S)
b_{20}	$-\frac{20}{30}, \frac{20}{30}$	$-1, 1$	$-1, 1$	$-\frac{40}{30}, \frac{40}{30}$	$-\frac{40}{30}, \frac{40}{30}$	$-1, 1$	$-1, 1$	$-\frac{20}{30}, \frac{20}{30}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_{15}	$-\frac{25}{30}, \frac{25}{30}$	$-1, 1$	$-1, 1$	$-\frac{35}{30}, \frac{35}{30}$	$-\frac{35}{30}, \frac{35}{30}$	$-1, 1$	$-1, 1$	$-\frac{25}{30}, \frac{25}{30}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\boldsymbol{b_{10}}$	$-\frac{30}{30}, \frac{30}{30}$	$\mathbf{-1, 1}$	$\mathbf{-1, 1}$	$-\frac{30}{30}, \frac{30}{30}$	$-\frac{30}{30}, \frac{30}{30}$	$\mathbf{-1, 1}$	$\mathbf{-1, 1}$	$-\frac{30}{30}, \frac{30}{30}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_5	$-\frac{35}{30}, \frac{35}{30}$	$-1, 1$	$-1, 1$	$-\frac{25}{30}, \frac{25}{30}$	$-\frac{25}{30}, \frac{25}{30}$	$-1, 1$	$-1, 1$	$-\frac{35}{30}, \frac{35}{30}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_0	$-\frac{40}{30}, \frac{40}{30}$	$-1, 1$	$-1, 1$	$-\frac{20}{30}, \frac{20}{30}$	$-\frac{20}{30}, \frac{20}{30}$	$-1, 1$	$-1, 1$	$-\frac{40}{30}, \frac{40}{30}$

Table 11 represents players' expected payoffs in the extended (malevolent) 3-color Ellsberg game, under the assumption of risk neutrality. The original pure Nash equilibria persist for the **K**eep-strategies. For the **S**witch-strategies, two mirrored Nash equilibria appear. As before, player D either chooses *bet-on-K* or *bet-on-U*, but does not *bet-on-C*. If the urn filler has a nonlinear utility function, only the Nash equilibria where player D bets on the known probability events are left (here Keep or Switch are irrelevant, since switching only alters the payoffs when subjects bet on an unkown event). As such, the predictions with respect to betting behavior are unchanged. The same holds true for the benevolent case.

Note that this is not surprising. As shown in Table 2, there is no way the experimenter can stack the urn in his favor in the 3-color urn. Therefore the suspicion argument has no bite in the 3-color Ellsberg game.

9 Appendix: Mixed Nash Equilibria

There are infinitely many Nash equilibria in mixed (non-pure) strategies. Denote by $\Delta(S_D)$ the set of all probability distributions on the set S_D. For any $\alpha = (\alpha_{1,3}, \alpha_{1,4}, \alpha_{2,3}, \alpha_{2,4})$ in $\Delta(S_D)$, $\alpha_{k,l}$ denotes the probability that player D's choices are (f_k, f_l) with $k \in \{1,2\}$ and $l \in \{3,4\}$. Denote by $\Delta(S_F)$ the set of all probability distributions on the set S_F. For any $\beta = (\beta_0, \ldots, \beta_n, \ldots, \beta_{20})$ in $\Delta(S_F)$, β_n is the probability that player F chooses a composition of the urn consisting of b_n black balls with $n = 0, \ldots, 20$. A pair of probability distributions (α^*, β^*) constitute a Nash equilibrium in mixed strategies for risk neutral players whenever they satisfy the following condition:

$$\begin{array}{lll} \alpha^* \in \Delta(S_D) & \text{such that} & \alpha_{1,3} = \alpha_{2,4} \quad \text{and} \quad \alpha_{1,4} = \alpha_{2,3}, \\ \beta^* \in \Delta(S_F) & \text{such that} & \sum_{n=0}^{20} \beta_n (\frac{20+b_n}{30}) = 1. \end{array} \tag{4}$$

The set of mixed equilbria is identical in the malevolent and the benevolent game.

References

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

Al-Najjar, N. I., and J. Weinstein (2009a): "The Ambiguity Aversion Literature: A Critical Assessment," *Economics and Philosophy*, 25, 249–284.

——— (2009b): "Rejoinder: The Ambiguity Aversion Literature: A Critical Assessment," *Economics and Philosophy*, 25, 357–369.

Alon, S., and D. Schmeidler (2014): "Purely Subjective Maxmin Expected Utility," *Journal of Economic Theory*, 152, 382–412.

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

Bade, S. (2011): "Ambiguous Act Equilibria," *Games and Economic Behavior*, 71, 246–260.

Camerer, C., and M. Weber (1992): "Recent Developments in Modelling Preferences: Uncertainty and Ambiguity," *Journal of Risk and Uncertainty*, 5, 325–370.

Casadesus-Masanell, R., P. Klibanoff, and E. Ozdenoren (2000a): "Maxmin Expected Utility over Savage Acts with a Set of Priors," *Journal of Economic Theory*, 92, 35–65.

——— (2000b): "Maxmin Expected Utility through Statewise Combination," *Economics Letters*, 66, 49–54.

Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2011): "Uncertainty Averse Preferences," *Journal of Economic Theory*, 146, 1275–1330.

Dominiak, A., P. Duersch, and J.-P. Lefort (2012): "A Dynamic Ellsberg Urn Experiment," *Games and Economic Behavior*, 75, 625–638.

Dominiak, A., and W. Schnedler (2011): "Attitudes Toward Uncertainty and Randomization: An Experimental Study," *Economic Theory*, 48, 289–312.

Dow, J., and S. R. D. C. Werlang (1994): "Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction," *Journal of Economic Theory*, 64, 305–324.

Eichberger, J., and D. Kelsey (2000): "Non-Additive Beliefs and Strategic Equilibria," *Games and Economic Behavior*, 30, 183–215.

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

——— (2001): *Risk, Ambiguity and Decision*. Routledge, New York, Published version of Ellsberg's 1962 doctoral thesis at Harvard University.

Ghirardato, P. (2002): "Revisiting Savage in a Conditional World," *Economic Theory*, 20, 83–92.

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

Ghirardato, P., F. Maccheroni, M. Marinacci, and M. Siniscalchi (2003): "A Subjective Spin on Roulette Wheels," *Econometrica*, 71, 1897–1908.

Gilboa, I., and M. Marinacci (2013): "Ambiguity and the Bayesian Paradigm," in *Advances in Economics and Econometrics: Theory and Applications, Tenth World Congress of the Econometric Society*, ed. by D. Acemoglu, M. Arellano, and E. Dekel. Cambridge University Press, New York.

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

Glaesmer, H., J. Hoyer, J. Klotsche, and P. Y. Herzberg (2008): "Die deutsche Version des Life-Orientation-Tests (LOT-R) zum dispositionellen Optimismus und Pessimismus," *Zeitschrift für Gesundheitspsychologie*, 16, 26–31.

Greiner, B. (2004): "An Online Recruitment System for Economic Experiments," in *Forschung und wissenschaftliches Rechnen 2003*, pp. 79–93. Ges. für Wiss. Datenverarbeitung, Göttingen.

Grèzes, J., C. Frith, and R. E. Passingham (2004): "Brain Mechanisms for Inferring Deceit in the Actions of Others," *Journal of Neuroscience*, 24, 5500–5505.

Hansen, L., and T. Sargent (2001): "Robust Control and Model Uncertainty," *American Economic Review*, 91, 60–66.

Hey, J., G. Lotito, and A. Maffioletti (2010): "The Descriptive and Predictive Adequacy of Theories of Decision Making under Uncertainty/Ambiguity," *Journal of Risk and Uncertainty*, 41, 81–111.

Hsu, M., M. Bhatt, R. Adolphs, D. Tranel, and C. F. Camerer (2005): "Neural Systems Responding to Degrees of Uncertainty in Human Decision-Making," *Science*, 310, 1680–1683.

Kadane, J. B. (1992): "Healthy Scepticism as an Expected-Utility Explanation of the Phenomena of Allais and Ellsberg," *Theory and Decision*, 32, 57–64.

Karni, E., F. Maccheroni, and M. Marinacci (2014): "Ambiguity and Non-expected Utility," in *Handbook of Game Theory*, ed. by P. Young, and S. Zamir, vol. 4. North Holland.

Kühberger, A., and J. Perner (2003): "The Role of Competition and Knowledge in the Ellsberg Task," *Journal of Behavioral Decision Making*, 16, 181–191.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005): "A Smooth Model of Decision Making under Ambiguity," *Econometrica*, 73, 1849–1892.

LEHRER, E. (2012): "Partially Specified Probabilities: Decisions and Games," *American Economic Journal: Microeconomics*, 4, 70–100.

LO, K. C. (1996): "Equilibrium in Beliefs under Uncertainty," *Journal of Economic Theory*, 71, 443–484.

MACCHERONI, F., M. MARINACCI, AND A. RUSTICHINI (2006): "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences," *Econometrica*, 74, 1447–1498.

MACHINA, M. J., AND M. SINISCALCHI (2014): "Ambiguity and Ambiguity Aversion," in *Handbook of the Economics of Risk and Uncertainty*, ed. by M. J. Machina, and W. K. Viscusi, vol. 1. North Holland.

MARINACCI, M. (2000): "Ambiguous Games," *Games and Economic Behavior*, 31, 191–219.

MUKERJI, S. (2009): "Foundations of Ambiguity and Economic Modeling," *Economics and Philosophy*, 25, 297–302.

OZDENOREN, E., AND J. PECK (2008): "Ambiguity Aversion, Games against Nature, and Dynamic Consistency," *Games and Economic Behavior*, 62, 106–115.

SAVAGE, L. J. (1954): *The Foundation of Statistics*. Willey, New York.

SCHMEIDLER, D. (1989): "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57, 571–587.

SCHNEEWEISS, H. (1968): "Spieltheoretische Analyse des Ellsberg-Paradoxons," *Zeitschrift für die gesamte Staatswissenschaft*, 124, 249–256.

——— (1973): "The Ellsberg Paradox From the Point of View of Game Theory," *Inference and Decision*, 1, 65–78.

Segal, U. (1987): "The Ellsberg Paradox and Risk Aversion: An Anticipated Utility Approach," *International Economic Review*, 28, 175–202.

Siniscalchi, M. (2009): "Two Out of Three Ain't Bad: A Comment on "The Ambiguity Aversion Literature: A Critical Assessment"," *Economics and Philosophy*, 25, 335–356.

Siniscalchi, M. (2011): "Dynamic Choice under Ambiguity," *Theoretical Economics*, 6, 379–421.

Trautmann, S., and G. van de Kuilen (2014): "Ambiguity Attitudes," in *Blackwell Handbook of Judgment and Decision Making*, ed. by G. Keren, and G. Wu. Blackwell.