On whom would I want to depend; Humans or nature? *

Mike Farjam†

We study in an experiment whether humans prefer to depend on decisions of other humans (social uncertainty) or states of nature (environmental uncertainty). In the social uncertainty treatments subjects depend only on past decisions of other humans. This is the first experiment that studies social uncertainty that does not derive from a strategic situation. The results indicate that even without any strategic context humans prefer lotteries where the distribution of outcomes is due to states of nature to lotteries where the distribution is due to decisions of humans. This holds even when distributions are identical and known to subjects.

JEL: C91, D81
Keywords: Ambiguity aversion, Experiment, Risk, Social Uncertainty

1. Introduction

Imagine you want to insure an expensive painting in your possession. The insurance company (which you trust) tells you the exact probability with which the picture will be stolen from you within 10 years. Furthermore, you learn that with the exact same probability the painting will be destroyed during some natural disaster. If you could only insure against one of the two events, which one would you choose? Would you be willing to pay more for one insurance policy than the other?

The experiment presented in this chapter studies whether humans prefer to depend on states of nature (*environmental uncertainty*) or other humans' decisions (*social uncertainty*). In the lottery choice experiment that we present in this chapter, subjects can choose the type of uncertainty on which outcomes of lotteries depend. Different from

*I thank the Max Planck Society for financial support through the International Max Planck Research School on Adapting Behavior in a Fundamentally Uncertain World. Special thanks also to Alexia Gaudeul, Oliver Kirchkamp, Anna Merkel and the participants of the 2014 IMPRS Uncertainty Summer School in Jena for their feedback and ideas on designing this experiment. I use R (2014) for the statistical analysis.

†University Jena, School of Economics, Carl-Zeiss-Str. 3, 07737 Jena, Mike.Farjam@uni-jena.de, ☎ +49 3641 930431, FAX +49 3641 943202.

previous experiments studying these types of uncertainty, the social uncertainty in the experiment does not derive from choices of others in a game or any kind of strategic context and the decisions that other humans took are not morally loaded (e.g. choosing the strategy in a trust game). This operationalization of social uncertainty allows for a clean comparison of preferences with regard to social versus environmental uncertainty.

A rational, expected value maximizing agent would not care whether the value it tries to maximize is affected by uncertainty due to some natural disaster or a human stealing from it. Humans, since they are not perfectly rational, may (e)valuate both types of situations very differently in terms of subjective probabilities and emotions involved. Abdellaoui et al. (2011) find that subjects prefer to depend on lotteries where the uncertainty of winning is determined by the weather, compared to a stock market. On a similar line McCabe, Rigdon, and Smith (2003) find that humans do not just care about the probabilities with which a decision is taken in a social context but also about the intentions that led to the decisions.

A clean comparison between environmental and social uncertainty is hard to achieve in the lab. Social uncertainty arises naturally in strategic interaction between two or more humans, e.g. in a game-like setting. This is why (to the best of our knowledge) in all experiments studying social uncertainty, the uncertainty derives from someones action in a strategic situation. Section 2.1 reviews these experiments and section 2.2 discusses recent findings in decision neuroscience that show potential neural mechanisms leading to different behavior under both kinds of uncertainty.

Carrubba, Yuen, and Zorn (2007) argue that strategic interaction – to some extend – can be seen as a lottery, where one depends on the actions of others without knowing which actions the others will take. Strategic uncertainty, however, differs in more than one aspect from environmental uncertainty. While environmental uncertainty involves only oneself depending on nature (a mechanism without intentions), strategic uncertainty includes interdependence between humans and includes beliefs about others intentions. Furthermore, the options that one can choose from create (often morally loaded) externalities on others. This in turn may lead to social preferences over outcomes. Dana, Weber, and Kuang (2007) show that in morally loaded contexts people actually prefer uncertainty. Summarized, one is not making a ceteris paribus comparison when comparing environmental and strategic uncertainty since in the latter, humans encounter more than just a decision problem. To overcome this flaw, we designed an experiment in which the only difference between a social and an environmental uncertainty condition is the source of uncertainty. The social uncertainty involved in treatments of the experiment will therefore not derive from strategic interaction.

We encounter situations of environmental uncertainty that do not arise from strategic interaction regularly. Examples are situations where we depend on decisions that others made without knowing their consequences, often we even depend on actions of others that are not the result of any conscious decision. As an example, an employee opening an email attachment that contains a computer virus may be absolutely unaware of causing harm to the company he is working for. For the employer the source of uncertainty, which he tries to protect his IT from, is not arising from a strategic decision of his employee, while it is also not environmental uncertainty because the source of uncertainty is human behavior.

Additionally to the differentiation between environmental and social uncertainty we also control the measurability of uncertainty (Knight, 1921). “Measurable” uncertainty (henceforth: *risk*) is characterized as a situation in which the probabilities with which all events occur are known. “Unmeasurable” uncertainty (henceforth: *ambiguity*[1]) is characterized by the absence of known probabilities. In his seminal thought experiment Ellsberg (1961) argues that humans generally avoid ambiguous lotteries in favor of lotteries were the distribution is known. Many experiments since have found that humans indeed are ambiguity averse (Camerer and Weber, 1992).

Table 1 shows the 4 possible combinations of the two dimensions along which uncertainty will be distinguished in this chapter. The conditions in the experiment (presented in section 3) correspond to the cells in the table.

The research questions of this chapter are:

1. Do humans prefer lotteries with social uncertainty to environmental uncertainty or vice versa?

2. Do humans have different risk preferences in lotteries where they depend on humans instate of nature?

3. Do we see ambiguity preferences in social uncertainty and are they the same as in environmental uncertainty?

	Distribution known	Distribution unknown
Social	Social Risk (SR)	Social Ambiguity (SA)
Environmental	Environmental Risk (ER)	Environmental Ambiguity (EA)

Table 1: Uncertainty concepts implemented

2. Literature

2.1. Experiments with strategic uncertainty

Heinemann, Nagel, and Ockenfels (2009) elicit certainty equivalents of two lotteries in an experiment. In the first lottery the probability distribution of winning and loosing depends on the actions of other players in a coordination-game (strategic uncertainty condition), in the second lottery the outcome depends on the roll of a die (environmental uncertainty condition). Heinemann, Nagel, and Ockenfels find that certainty equivalents of both lotteries are similar. From a ceteris paribus point of view the comparison Heinemann, Nagel, and Ockenfels are making is problematic since their environmental

[1]Note that in some parts of the literature this is referred to as *uncertainty*. Unfortunately there is no consistency with using these terms in the literature.

uncertainty condition is characterized by known probabilities and thus is a risky decision problem, while the strategic uncertainty condition is characterized by ambiguity. Furthermore, all arguments made in section 1 when it comes to comparing a game with a decision problem apply.

Bohnet and Zeckhauser (2004) – and similarly Bohnet et al. (2008) – addressed many of the issues when it comes to a clean comparison between risk and strategic/social uncertainty. In their experiment subjects play a binary trust game, where half of the subjects are first-movers (Trustors) and the other half second-movers (Trustees). They ask first-movers for the minimum acceptable probability (MAP) of playing against a trustworthy second-mover for which they would be willing to make a trust move. First- and second-movers than are randomly matched and first-movers will play the trust move if their MAP is smaller than the fraction of trustworthy first movers. In asking for the MAP Bohnet and Zeckhauser transform the strategic interaction to a risky decision problem. They find that MAPs of first-movers are generally higher when the second-mover's move is based on a decision compared to when the strategy was determined by a random device. The fact that subjects ask for an extra risk premium when depending on humans' choices compared to depending on a mechanism is interpreted as evidence for "betrayal aversion"[2]. We argue that Bohnet and Zeckhauser jump to conclusions since as in Heinemann, Nagel, and Ockenfels (2009) they compare a situation where one depends on a neutral mechanism without any intentions with a situation where one depends on a human who had (morally loaded) intentions when taking the conscious decision. An alternative explanation for the finding of Bohnet and Zeckhauser could be inequality aversion (as discussed by Bolton and Ockenfels (2010)) or what one could refer to as "human aversion" (humans just do not like to depend on other humans).

Similar to Bohnet and Zeckhauser (2004), Fairley et al. (2014) let subjects play a trust game as both Trustee and Trustor. When in the role of a Trustor, subjects were assigned to four random Trustees and one random decision of the Trustees was played. There were thus five potential probabilities of encountering a trustworthy Trustee depending on the number of trustworthy Trustees. Trustors had to indicate the amount for each potential probability that they would be willing to pay in order to play the lottery. Fairley et al. compare the amounts per probability with amounts that had to be indicated for standard risky lotteries with the same 5 potential probabilities of winning. They find that choices in lotteries and social lotteries are not correlated.

The experiments discussed do not give a clear picture on how the various operationalizations of strategic and social uncertainty on the one hand and environmental uncertainty on the other relate to each other. However, since all results were obtained by comparing environmental uncertainty with uncertainty based on some strategic action of another human, it remains unclear how these results relate to a comparison without a strategic context (as we will do in this chapter).

[2]This paper is the core reference in chapter **??** and is discussed in more detail there.

2.2. Decision Neuroscience

Experiments in the area of decision neuroscience and neuroeconomics indicate that subjects react differently to situations with social and environmental ambiguity simply because the one involves humans and this leads to activation of different brain areas. Lauharatanahirun, Christopoulos, and King-Casas (2012) study differences in brain activity (via fMRI) when subjects engage in a Trust Game compared to a lottery where the outcome depends on a random mechanism. They find that especially the amygdala (related to emotions) is more active in the Trust Game.

Sanfey (2007) argues that emotions are crucial in economic decisions that involve other humans and brain areas that involve emotion regulation are therefore more active in situations with social uncertainty. In Sanfey et al. (2003) the authors find that the anterior insula is more active when subjects get unfair offers in an Ultimatum Game where the proposer is a human compared to the proposer being a random device. The latter seems to evoke less emotions than interaction with humans. This finding is robust across different types of games like Rock-Paper-Scissors (Chaminade et al., 2012), prisoners dilemma game (Krach et al., 2008; Rilling et al., 2004) and trust games (McCabe et al., 2001). These experiments also show that humans invest more effort when their counterpart is human.

McCabe et al. (2001) and Rilling et al. (2004) similarly find that brain areas known to be crucial for theory of mind (the ability to put oneself into the shoes of someone else) are more active when subjects know they are playing with humans than when they know they are playing with computers. On the same line Rilling et al. (2002) study activation in the striatum, which is a crucial brain area when social decisions have to be made, when subjects interact with computers and robots. This area is less active when dealing with situations that involve environmental uncertainty.

Although the neuroeconomic literature shows that from a perceptional point of view environmental and social uncertainty are different, the literature does not allow for a clear prediction whether humans dislike or prefer situations with social uncertainty to environmental uncertainty.

3. Methods

3.1. General design

Each subject was confronted with four lotteries in a random order. Each of these lotteries contained one of the kinds of uncertainty presented in table 1. After each lottery subjects indicated (through a multiple price list) how high a compensation has to be in order for them to renounce the uncertain payoff of the lottery. After the four lotteries subjects had to play one of the previous lotteries an additional time and could indicate which one they want this to be. The entire experiment was computerized using z-Tree (Fischbacher, 2007).

3.2. Treatments

In the lotteries in treatment ER (Environmental Risk) and EA (Environmental Ambiguity) subjects had to guess the result of a card draw. Cards were drawn from a stack of 10 cards. Each card was either an *A*- or a *B*-card. In ER the distribution was 50/50 and subjects were explicitly told so. In EA subjects did not know the distribution and were told that any number of *A*-cards from 0–10 was possible and that the remaining of the 10 cards were B-cards. A random number would be generated by the computer determining one of the 10 cards. Subjects had to guess if it was an *A*- or a *B*-card. If their guess was correct they would receive 100 ECU for that task otherwise 0. 10 ECU were worth 1 Euro.

The lotteries in treatments SR (Social Risk) and SA (Social Ambiguity) were almost identical to ER and EA. Before the experiment took place a survey was conducted among students that did not participate in the experiment. In treatments SR and SA instead of cards, the computer would draw a random answer from the survey. Subjects in the experiment were told that 10 participants of the survey had to indicate their preference towards two different pictures. Subjects got no information on the content of pictures or the artists. They were only told that the one picture would be referred to as picture *A* and the other as *B*. The task of the subject was to guess whether the participant of the survey drawn by the computer preferred picture *A* or *B*. In SR subjects knew that the distribution was 50/50 and in SA the distribution was unknown to the subject. It was also made clear that the participants in SR are not the same participants as in SA.[3]

3.3. Dependent variables

To elicit subjects' preferences regarding the lotteries in each treatment two different elicitation mechanisms were used. After subjects played all the lotteries in all 4 treatments they were presented with a table in which the rows contained all possible pairwise comparisons between the 4 treatments. Per pair of treatments subjects indicated which of the treatments they would prefer to play as an additional 5th treatment. One of these pairs of treatments was chosen randomly by the computer and the subject's preferred treatment was then played. To rule out order effects with regard to the order in which treatments are paired, the comparison that subjects had to make within each pair was not directly about treatments, but about rounds. Since the treatment per round was determined randomly, the treatments in each row of the comparison table also have a random order per subject.

As a second elicitation mechanism subjects had to fill in a multiple price list in every treatment after taking their guess on *A* or *B* in the lottery. In this multiple price list subjects had to choose 10 times between a secure payoff that ranges from 9 to 99 ECU in steps of 10 ECU and an uncertain payoff which would be either 0 or 100 ECU depending on the result of the lottery. As in Holt and Laury (2002) the relative frequency of lottery-choices will be used as a measure for preference for that lottery.

[3]This surveys of course really took place, done with pictures from Klee and Kandinsky on the campus of the university in Jena.

3.4. Instructions and controls

Before subjects got instructions per treatment they watched a video (2:36 minutes) with general instructions. The video (with English subtitles) is available at `https://youtu.be/Qtn1q-4jB7M`. Appendix A.1 contains the full set of instructions describing the lotteries in each treatment.

To avoid that differences in preferences regarding the treatments may be driven by complexity aversion (Sonsino, Benzion, and Mador, 2002), the instructions per treatment were as similar as possible in terms of length, wording and syntactic structure. Table 2 gives a ruff overview of the complexity of the instructions in each treatment.

To avoid order effects the treatments/lotteries were presented in a random order per subject. To control whether subjects understood the distribution of cards in each treatment they were asked at the end of the experiment how many *A*- and *B*-cards/survey answers they expects to be in the stack/survey per treatment. The answers were incentified for correctness.

Treatment	Sentences	Words	Characters
Environmental Risk	4	51	272
Environmental Ambiguity	4	54	288
Social Risk	4	47	288
Social Ambiguity	4	51	301

Table 2: Comparing the complexity of the (German) instructions per treatment

4. Results

4.1. Subjects

88 subjects participated in this study. All subjects were recruited via ORSEE (Greiner, 2004). Since studies like Dohmen et al. (2011) show that risk-preferences differ between genders, only male subjects were recruited to reduce within sample variability. All sessions were run in May 2015 in the laboratory of the Friedrich Schiller University Jena. Most of our subjects were students.

4.2. Payoff

To avoid endowment effects only one of the treatments was chosen randomly at the end of the session for payoff. Subjects earned on average 7 Euro (including a 2.5 Euro show-up fee). The experiment (including the payment) took about 30 minutes.

4.3. Pairwise comparison

This subsection compares how often subjects preferred one treatment above one of the other 3 treatments. Thus, the maximum number of times a lottery could be preferred above other lotteries is 3 and the minimum is 0. Table 3 shows the result of a pairwise comparison of all 4 treatments.

	ER	EA	SR
EA	0.0000		
SR	0.0018	0.0000	
SA	0.0000	0.5059	0.0000

p-values adjusted for multiple testing with the Holm–Bonferroni method.

Table 3: P-values of a pairwise comparison between treatments with a Wilcox rank-sum test

From table 3 and figure 1 (left) we can derive the following order of preferences: $ER \succ SR \succ EA \sim SA$. The fact that $ER \succ EA$ and $SR \succ SA$ indicates that subjects were ambiguity averse. Ambiguity aversion thus seems to be found not just in the environmental but also in the social domain. Furthermore, 66 % of subjects preferred environmental risk above social risk in the direct comparison.

To check whether the results are driven by an order effect (despite the randomization of the order of treatments per subject) two mixed effects model were used to predict the preferences of subjects regarding a treatment. An ANOVA comparing the model with and without the order of treatments as a fixed factor shows that the model without order performs better in terms of the AIC and BIC. The preferences subjects had regarding the treatments seem thus not to be affected by the order in which the treatments were presented.

4.4. Multiple price list

Subjects had to make 10 choices between a fixed secure payoff and the uncertain payoff of the lottery. From the subject's point of view the expected payoff of all lotteries was 50 ECU. Hence rational, risk-neutral agents would choose the payoff of the lottery if the fixed payoff is below 50 ECU and choose the fixed payoff if it is above 50 ECU. Given the multiple price list subjects were confronted with (see appendix A.3) one would expect that subjects choose 5 times the fixed payoff and 5 times the payoff of the treatment.

On average subjects chose the fixed payoff 5.6 times, which is more often than a rational, risk-neutral agent would and which indicates that subjects were (as usually found) risk-averse. 33 % of subjects chose the fixed lottery equally often in all treatments, indicating that they were indifferent between all treatments. Figure 1 (right) shows that the distributions of choices for the payoff of the lottery per treatment are almost identical. Furthermore, 12.5 % of subjects had more than one switching point in the multiple price list and 10 % of subjects had more than 1 switching point which indicates that subjects did not understand how a multiple price list works. Multiple switching points are often found when using multiple price lists for preference elicitation (Andersen et al., 2006).[4]

[4] Appendix A.2 shows the results of a pilot study of this experiment where a BDM mechanism (Becker, DeGroot, and Marschak, 1964) was used to elicit preferences. Results obtained by the BDM mechanism are almost the same as with the multiple price list.

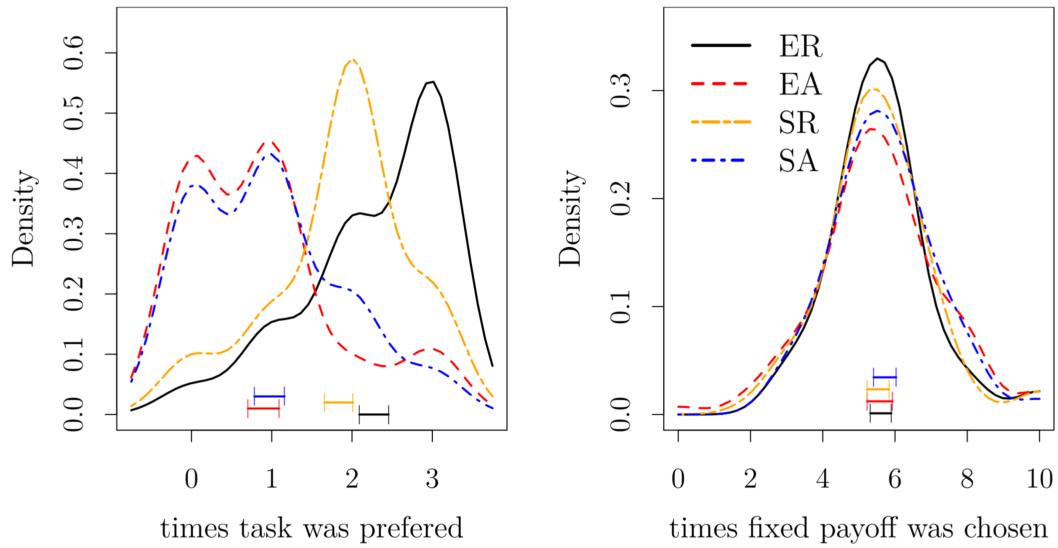

Per distribution of answers the 95% bootstrap confidence interval of the mean is shown.

Figure 1: Left, distributions on how often subjects preferred each treatment above the other 3 treatments. Right, distributions of choices for the lottery instead of the safe payment per treatment in the multiple price list.

Summarized, analysis of the data with regard to choices in the multiple price list seems problematic. The statistical analysis of subjects' choices in the multiple price list will therefore only be rudimentary.

A Friedman-test comparing the mean number of times the lottery was preferred above the fixed amount shows that the differences between treatments with regard to choices in the multiple price list is not significant (p = 0.37, $\chi^2 = 3.16$). A Bartlett-test comparing the variance between groups finds that the only significant difference regarding variance can be found between treatments ER and EA (not adjusted for multiple testing: p = 0.029, $K^2 = 4.77$). Neither of the other comparisons are significant below the 0.1 level. Looking at figure 1 (right) it seems unlikely that differences would become significant with a larger sample. It thus seems that although the previous subsection shows that subjects prefer some treatments above others, the treatment they are in does not change their risk-preference regarding an opt-out option from that treatment.

4.5. Robustness checks

To check whether the results may be driven by different believes about the distributions in each treatment subjects had to guess the number of A's in each treatment. Figure 2 (left) shows that on average subjects had the same expectations in all treatments and that the great majority in ER and SR understood that there were exactly 5 A's in these treatments.

Furthermore, we measured per subject how long they remained on screens that presented the description of a treatment (the exact instructions can be found in appendix A.1). The time spend on this screen can be seen as a proxy for perceived complexity of the description. Figure 2 (right) shows that there were hardly any differences between treatments EA and SA and treatments ER and SR. However, subjects remained on average 7.2 seconds longer on the screens that presented instructions in treatment EA than in treatment ER and 6.2 seconds longer on screens of SA than SR.

Summarized, the fact that subjects preferred ER to SR seems not to be driven by a misunderstanding of the two treatments or differences in perceived complexity. However, the fact that subjects generally prefer SR to SA and ER to EA may be explained by both: ambiguity and complexity aversion.

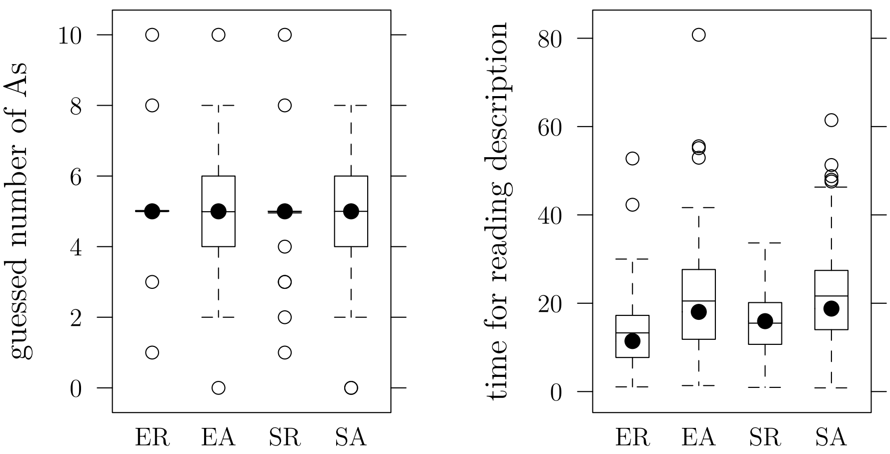

Box-and-whisker plots per treatment. Dots represent medians, lines means.

Figure 2: Left, subjects' believe about the frequency of A's. Right, time spend for reading description on distribution of A's

5. Discussion

In the experiment presented, subjects had the choice between engaging in a lottery where they depend on a distribution of outcomes based on decisions of humans and a lottery where the distribution does not depend on humans but only on nature. When the distributions of outcomes was ambiguous, preferences with regard to both lotteries were on average the same. In the risky lotteries where the distribution of outcomes was known and identical subjects clearly preferred to depend not on other humans but on nature.

This finding is in line with the conclusions from the experiment by Bohnet and Zeckhauser (2004). While Bohnet and Zeckhauser show that humans demand a "risk premium" in order to accept strategic dependence on another human compared to a random device, we show that humans demand this premium even outside any strategic context. Bohnet and Zeckhauser conclude that subjects are betrayal averse. In the light of our results Bohnet and Zeckhauser's finding may not at all be related to an aversion towards betrayal but simply to the fact that one has an aversion to depend on others (irrespective of the others intention).

To the best of our knowledge this experiment shows for the first time that humans care about the fact that they depend on other humans irrespective of the outcome and even in the absence of intentions of the human on whom they depend. This suggests that differences in perception between social and environmental uncertainty are more fundamental than expected. This finding is in line with the claim of McCabe, Rigdon, and Smith (2003) that humans do not just care about the probabilities with which an event occurs but also about the process that led to these probabilities. In our experiment we find that even when the distribution on which an outcome in a lottery depends is identical and known, humans may systematically prefer one lottery to the other.

In our experiment subjects had clearly different preferences regarding environmental and social risk, yet they did not have different risk preferences once facing one of the two types of risk in a lottery. This may mean that humans given the choice whether they want to depend on humans or nature would chose for nature, but that humans once they are in a specific situation do not differ in terms of behavior in both conditions. However, risk preferences are not the only indicator for behavior and future studies have to look more carefully at how human behavior under environmental and social risk differs.

We can only speculate on the reasons why subjects preferred to depend on nature and not on other humans' decisions. Besides the differences in the neural mechanism that may work depending on the kind of risk (see section 2.2), Heath and Tversky (1991) show that humans generally prefer to engage in lotteries which they know more about or feel more knowledgeable. This holds even if the knowledge is not relevant for the outcome of the lottery. Subjects in the experiment may have felt less knowledgeable about lotteries with social uncertainty compared to environmental uncertainty, simply because being confronted with decisions of other humans was perceived as less transparent by subjects than being confronted with the letter on a card. This is of course only a speculation and the cognitive mechanism that makes humans prefer environmental above social risk may be a different one.

The results we obtained also have methodological implications. One method that is used to control for strategic uncertainty in a game is to have a baseline condition with the normal game and a treatment condition with a computerized player where the algorithm of the computer is known to subjects (e.g. Koch and Penczynski (2015)). Our results however show that playing against a computer may be perceived as fundamentally different from playing and depending on a human, irrespective of strategic considerations. Conclusions based on such a comparison are therefore problematic.

If we interpret strategic interaction as a lottery where the outcome depends on a second player the results of this chapter suggest that the nature of the second player maybe an important feature in itself, irrespective of any strategic considerations. This suggests that humans may behave differently in strategic interaction with humans from strategic interaction with e.g. machines, even if they know that machine and human will behave identically. We live in a time where the number of interactions between humans and machines grows inexorably and more and more human-human interaction is replaced by human-machine interaction. It is therefore of great importance to study exactly how humans perceive strategic interaction with machines differently from interaction with humans and how this changes their (economic) decisions. Research with this objective can help in the design of e.g. expert systems whose advice will be accepted by humans and identify situations where human-machine interaction can be beneficial/harmful compared to human-human interaction.

References

Abdellaoui, Mohammed, Aurélien Baillon, Laetitia Placido, and Peter P Wakker (2011). "The rich domain of uncertainty: Source functions and their experimental implementation". In: *American Economic Review* 101.2, pp. 695–723.

Andersen, Steffen, Glenn W Harrison, Morten Igel Lau, and E Elisabet Rutström (2006). "Elicitation using multiple price list formats". In: *Experimental Economics* 9.4, pp. 383–405.

Becker, Gordon M, Morris H DeGroot, and Jacob Marschak (1964). "Measuring utility by a single-response sequential method". In: *Behavioral Science* 9.3, pp. 226–232.

Bohnet, Iris and Richard Zeckhauser (2004). "Trust, risk and betrayal". In: *Journal of Economic Behavior & Organization* 55.4, pp. 467–484.

Bohnet, Iris, Fiona Greig, Benedikt Herrmann, and Richard Zeckhauser (2008). "Betrayal aversion: Evidence from Brazil, China, Oman, Switzerland, Turkey, and the United States". In: *American Economic Review* 98.1, pp. 294–310.

Bolton, Gary E and Axel Ockenfels (2010). "Betrayal aversion: Evidence from brazil, china, oman, switzerland, turkey, and the united states: Comment". In: *American Economic Review* 100.1, pp. 628–633.

Camerer, Colin and Martin Weber (1992). "Recent developments in modeling preferences: Uncertainty and ambiguity". In: *Journal of Risk and Uncertainty* 5.4, pp. 325–370.

Carrubba, Clifford J, Amy Yuen, and Christopher Zorn (2007). "In defense of comparative statics: Specifying empirical tests of models of strategic interaction". In: *Political Analysis* 15.4, pp. 465–482.

Chaminade, Thierry et al. (2012). "How do we think machines think? An fMRI study of alleged competition with an artificial intelligence". In: *Frontiers in Human Neuroscience* 6.103.

Dana, Jason, Roberto A Weber, and Jason Xi Kuang (2007). "Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness". In: *Economic Theory* 33.1, pp. 67–80.

Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G Wagner (2011). "Individual risk attitudes: Measurement, determinants, and behavioral consequences". In: *Journal of the European Economic Association* 9.3, pp. 522–550.

Ellsberg, Daniel (1961). “Risk, Ambiguity, and the Savage Axioms”. In: *The Quarterly Journal of Economics* 75.4, pp. 643–669.

Fairley, Kim, Alan G Sanfey, Jana Vyrastekova, and Utz Weitzel (2014). *Trust and risk revisited.* Tech. rep. 59900. MPRA Paper.

Fischbacher, Urs (2007). “z-Tree: Zurich toolbox for ready-made economic experiments”. In: *Experimental Economics* 10.2, pp. 171–178.

Greiner, Ben (2004). “An Online Recruitment System for Economic Experiments”. In: *Forschung und wissenschaftliches Rechnen.* Ed. by Kurt Kremer and Volker Macho. Vol. 63, GWDG Bericht. Ges. für Wiss. Datenverarbeitung. Göttingen, pp. 79–93.

Heath, Chip and Amos Tversky (1991). “Preference and belief: Ambiguity and competence in choice under uncertainty”. In: *Journal of Risk and Uncertainty* 4.1, pp. 5–28.

Heinemann, Frank, Rosemarie Nagel, and Peter Ockenfels (2009). “Measuring strategic uncertainty in coordination games”. In: *The Review of Economic Studies* 76.1, pp. 181–221.

Holt, Charles A and Susan K Laury (2002). “Risk aversion and incentive effects”. In: *American Economic Review* 92.5, pp. 1644–1655.

Knight, Frank H (1921). *Risk, uncertainty and profit.* New York: Hart, Schaffner and Marx.

Koch, Christian and Stefan Penczynski (2015). *The Winner’s Curse: Conditional Reasoning & Belief Formation.* Tech. rep. 15-03. University of Mannheim / Department of Economics Working Paper Series.

Krach, Sören, Frank Hegel, Britta Wrede, Gerhard Sagerer, Ferdinand Binkofski, and Tilo Kircher (2008). “Can machines think? Interaction and perspective taking with robots investigated via fMRI”. In: *PLoS One* 3.7, e2597.

Lauharatanahirun, Nina, George I Christopoulos, and Brooks King-Casas (2012). “Neural computations underlying social risk sensitivity”. In: *Frontiers in Human Neuroscience* 6.213.

McCabe, Kevin, Daniel Houser, Lee Ryan, Vernon Smith, and Theodore Trouard (2001). “A functional imaging study of cooperation in two-person reciprocal exchange”. In: *Proceedings of the National Academy of Sciences* 98.20, pp. 11832–11835.

McCabe, Kevin A, Mary L Rigdon, and Vernon L Smith (2003). “Positive reciprocity and intentions in trust games”. In: *Journal of Economic Behavior & Organization* 52.2, pp. 267–275.

R Development Core Team (2014). *R: A Language and Environment for Statistical Computing.* ISBN 3-900051-07-0. R Foundation for Statistical Computing. Vienna, Austria.

Rilling, James K, David A Gutman, Thorsten R Zeh, Giuseppe Pagnoni, Gregory S Berns, and Clinton D Kilts (2002). “A neural basis for social cooperation”. In: *Neuron* 35.2, pp. 395–405.

Rilling, James K, Alan G Sanfey, Jessica A Aronson, Leigh E Nystrom, and Jonathan D Cohen (2004). “The neural correlates of theory of mind within interpersonal interactions”. In: *Neuroimage* 22.4, pp. 1694–1703.

Sanfey, Alan G (2007). “Social decision-making: insights from game theory and neuroscience”. In: *Science* 318.5850, pp. 598–602.

Sanfey, Alan G, James K Rilling, Jessica A Aronson, Leigh E Nystrom, and Jonathan D Cohen (2003). "The neural basis of economic decision-making in the ultimatum game". In: *Science* 300.5626, pp. 1755–1758.

Sonsino, Doron, Uri Benzion, and Galit Mador (2002). "The Complexity Effects on Choice with Uncertainty–Experimental Evidence". In: *The Economic Journal* 112.482, pp. 936–965.

A. Appendix

A.1. Instructions

This subsection contains the instructions in each treatment translated to English. To get an idea of the general length of the original instructions the German versions are printed in italics for the first four parts of the experiment.

A.1.1. Environmental Risk

We took 10 cards from a large stack of cards and shuffled these cards. On 5 of these cards an *A* is written, on the other 5 cards a *B* is written.
The computer will draw one of the 10 cards randomly. Guess whether there will be an *A* or *B* on the card.

Wir haben aus einem großen Stapel mit Karten 10 Karten genommen und diese zu einem Stapel gemischt. Auf 5 dieser Karten steht A*, auf den anderen 5 Karten steht ein* B*.*
Der Computer wählt eine der 10 Karten zufällig aus. Raten Sie ob auf der ausgewählten Karte A *oder* B *steht.*

A.1.2. Envirionmental Ambiguity

We took 10 cards from a large stack of cards and shuffled these cards. On 0 to 10 cards (the exact number is unknown to you) an *A* is written, on the other ones a *B* is written. The computer will draw one of the 10 cards randomly. Guess whether there will be an *A* or *B* on the card.

Wir haben aus einem großen Stapel mit Karten 10 Karten genommen und diese zu einem Stapel gemischt. Auf 0 bis 10 Karten (die genaue Zahl wissen Sie nicht) steht A*, auf den anderen* B*.*
Der Computer wählt eine der 10 Karten zufällig aus. Raten Sie ob auf der ausgewählten Karte A *oder* B *steht.*

A.1.3. Social Risk

We conducted a survey on the campus among 10 people and showed them two different pictures. 5 of the participants preferred picture *A*, the other 5 participants preferred picture *B*.
The computer will draw one random participant from the 10. Guess whether the participant drawn preferred picture *A* or *B*.

Wir haben in einer Befragung auf dem Campus 10 Menschen zwei verschiedene Bilder gezeigt. 5 der Befragten bevorzugten Bild A*, die anderen 5 Befragten bevorzugten Bild* B*.*

Der Computer wählt einen der 10 Befragten zufällig aus. Raten Sie ob der ausgewählte Befragte Bild A *oder* B *bevorzugte.*

A.1.4. Social Ambiguity

We conducted a survey on the campus among 10 people and showed them two different pictures. 0 to 10 participants (the exact number is unknown to you) of the participants preferred picture *A*, the others preferred picture *B*.
The computer will draw one random participant from the 10. Guess whether the participant drawn preferred picture *A* or *B*.

Wir haben in einer Befragung auf dem Campus 10 Menschen zwei verschiedene Bilder gezeigt. 0 bis 10 Befragte (die genaue Zahl wissen Sie nicht) bevorzugten Bild A*, die anderen Bild* B*.*
Der Computer wählt einen der 10 Befragten zufällig aus. Raten Sie ob der ausgewählte Befragte Bild A *oder* B *bevorzugte.*

A.1.5. Choice Task 5

In the following table 2 of the 4 parts of the experiment are put next to each other. Indicate per row which of the parts of the experiment you prefer. The computer will randomly chose one of the rows. The part of the experiment that you preferred in the determined row will be played as an additional part of this experiment.[5]

Option A	Option B
1st part of the experiment	2nd part of the experiment
1st part of the experiment	3rd part of the experiment
1st part of the experiment	4th part of the experiment
2nd part of the experiment	3rd part of the experiment
2nd part of the experiment	4th part of the experiment
3rd part of the experiment	4th part of the experiment

A.2. BDM elicitation mechanism

In a pilot experiment the BDM mechanism was used (Becker, DeGroot, and Marschak, 1964) to elicit subjects' (N = 71) preferences with regard to the treatments. Figure 3 compares the results obtained with the help of the BDM mechanism with those from the experiment discussed in the paper. The distributions of preferences hardly differ between those two elicitation mechanisms.

[5]Additionally subjects got a description of each part of the experiment as shown in the previous subsections.

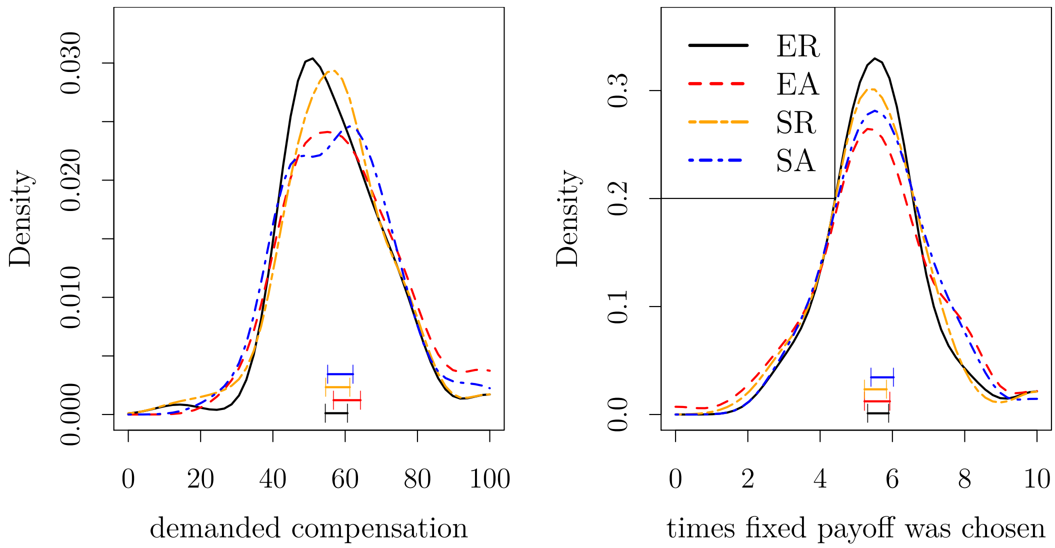

Figure 3: Left, preferences w.r.t. the treatments elicited by the BDM-mechanism; Right, elicited with the help of a multiple price list. Per treatment the 95% bootstrap confidence interval of the mean is shown.

A.3. Multiple Price List

The multiple price list presented to subjects after each treatment:

Option A	Option B
Get 9 ECU for sure	your payoff depends on the lottery
Get 19 ECU for sure	your payoff depends on the lottery
...	...
Get 89 ECU for sure	your payoff depends on the lottery
Get 99 ECU for sure	your payoff depends on the lottery