

Utrecht University
Faculty of Geosciences

**Understanding Climate Change:
Economic Behaviour under Risk and Uncertainty**

Master Thesis

by

Cristina-Ioana Apetrei

Supervisors: Dr. Carel Dieperink & Prof. dr. Bert de Vries

February 2010

Utrecht University
Faculty of Geosciences

**Understanding Climate Change:
Economic Behaviour under Risk and Uncertainty**

C.I. Apetrei: 3206564

Thesis submitted in partial fulfilment of the requirements for the degree of
Master of Science in Sustainable Development

Supervisors: Dr. Carel Dieperink & Prof. dr. Bert de Vries
45 ECTS

February 2010

Acknowledgements

Solving the climate change issue is no easy task. Nor is it to try to make the world a better place. Nor is it to sustain one's efforts for many months in an attempt to complete a meaningful project.

But, after more than two years at Utrecht University, I can tell that any of the above is worthwhile doing when you are surrounded by wonderful and supportive people.

I would like to thank my supervisors, prof. Bert de Vries and prof. Carel Dieperink, for having guided me along this road and, even most importantly, for continuously believing in me and in my interdisciplinary approach of the research topic.

There are few words to express my thankfulness to prof. Vincent Buskens, whose gratuitous help and support during the experiments were invaluable.

I am also grateful for his comments, advices and encouragements to Markus Brede, whom I've only met a few times, but in what I thought to be key moments of my research.

To Christian Michel, who developed the interactive platform that we used in our experiments, I would like to thank for being such a perfect team-mate.

Last, but not least, I am forever grateful to my family and friends who have patiently been by my side even in the darkest of nights of what I sometimes thought to be a never-ending story.

Now, when my journey is complete, I am confident that with the support of such people anything can be accomplished, even solving the climate change issue.

Contents

LIST OF FIGURES	IX
LIST OF TABLES	X
LIST OF INSERTS	X
ABBREVIATIONS	XI
1. INTRODUCTION.....	1
1.1. CLIMATE CHANGE AS AN OBSERVABLE PHENOMENON	1
1.2. CLIMATE CHANGE AS A POLICY PROBLEM	2
1.3. RESEARCH OBJECTIVE AND MAIN QUESTION	3
1.4. RESEARCH FRAMEWORK	4
2. DECISION THEORY	7
2.1. THE PHILOSOPHY OF DECISION MAKING	7
2.2. FIELD OVERVIEW & STRUCTURE OF THE CHAPTER	8
2.3. AN X-RAY OF A DECISION SITUATION	9
2.4. TRADITIONAL DECISION THEORY	12
2.4.1. <i>Delineating the Field</i>	12
2.4.2. <i>Dealing with Uncertainty: Probabilities of Events</i>	13
2.4.3. <i>Dealing with Outcomes: Expectations, Payoffs, Utilities</i>	20
2.4.4. <i>Rational Choice Theory</i>	23
2.5. ALTERNATIVE THEORIES	25
2.5.1. <i>Behavioural Decision Theory: New Developments</i>	25
2.5.2. <i>Gauthier & Rationality of Commitments</i>	27
2.6. CONCLUSIONS	27
3. GAME THEORY	29
3.1. GAME THEORY AS INTERACTIONIST DECISION MAKING	29
3.2. TRADITIONAL GAME THEORY	29
3.2.1. <i>Key Concepts</i>	30
3.2.2. <i>Game Theory and Decision Making under Complete Ignorance</i>	35
3.2.3. <i>Prisoner's Dilemma Game</i>	38
3.2.4. <i>Advanced Game Theory</i>	39
3.3. BEHAVIOURAL GAME THEORY	39
3.3.1. <i>Methodological Considerations</i>	40
3.3.2. <i>A Few Empirical Findings</i>	40
3.4. CONCLUSIONS	41
4. ENVIRONMENTAL COMMONS DILEMMA AND CLIMATE CHANGE.....	43
4.1. INTRODUCING THE CONCEPTS	43
4.2. CLIMATE CHANGE PARTICULARITIES	45
4.3. BEHAVIOUR IN ENVIRONMENTAL COMMONS DILEMMAS	45
4.3.1. <i>The Prisoner's Dilemma as a Model: Uses and Limitations</i>	45
4.3.2. <i>Cooperation: Explanatory Models</i>	47
4.3.3. <i>Protection Motivation Theory</i>	49
4.4. LESSONS FROM PREVIOUS CLIMATE CHANGE EXPERIMENTS	52
4.5. CONCLUSIONS	54
5. MODELLING THE CLIMATE CHANGE DECISION CONTEXT.....	55
5.1. A VIEW ON THE WORLD	55
5.2. CONCEPTUAL MODEL	56
5.2.1. <i>Variables</i>	58
5.2.2. <i>Hypotheses</i>	60
6. TESTING THE MODEL.....	61
6.1. METHOD	61
6.1.1. <i>Why The Need of a Game?</i>	61

6.1.2. CLIMEX – An Interactive Decision-Making Environment	62
6.1.3. Main Game Setup Description.....	62
6.2. FORMALIZING THE MODEL	63
6.2.1. Dependent variable: Individual investment behaviour	63
6.2.2. Independent variables.....	64
6.3. EXPERIMENTAL DESIGN	66
6.3.1. Treatments	66
6.3.2. Participants, Format and Questionnaire	68
6.3.3. Game Parameterization.....	68
6.3.4. Outputs: The Data	70
7. EXPERIMENT OUTCOMES.....	71
7.1. DATA ANALYSIS.....	71
7.1.1. Measures of the Dependent Variable.....	71
7.1.2. Risk Attitude Profiles	73
7.1.3. Other Covariates	77
7.1.4. Other Indicators: Income Effects.....	78
7.2. RESULTS.....	79
7.2.1. Rationality and Malignancy.....	79
7.2.2. Rationality and Uncertainty.....	82
7.2.3. Rationality and Order.....	83
7.2.4. Effects on Investment Level.....	85
7.2.5. Characteristics of the Participants	87
8. CONCLUSIONS AND RECOMMENDATIONS	89
REFERENCES	91
APPENDICES	97
APPENDIX A. CLIMEX	99
APPENDIX B. EXPERIMENT INSTRUCTIONS.....	105
APPENDIX C. RISK MODES	115
APPENDIX D. ADDITIONAL INSTRUCTIONS SHEET.....	117
APPENDIX E. QUESTIONNAIRE	119
APPENDIX F. OPTIMUM INVESTMENTS	125
APPENDIX G. RANDOM VARIABLES	127
APPENDIX H. STATISTICAL TABLES	129

List of Figures

Figure 1. Research Framework	5
Figure 2. Decision Matrix	10
Figure 3. Paul's Decision Matrix	12
Figure 4. "Six Face Up" Bet - Payoff Matrix	21
Figure 5. Coordination Game: Payoff Matrix	30
Figure 6. Cooperation Game in Extensive Form	31
Figure 7. Payoff Matrix with Best Replies for Paul and Ana	33
Figure 8. Modified Payoff Matrix	35
Figure 9. Paul's Decision: Payoff Matrix	36
Figure 10. Hypothetical Payoff Matrix with Three Alternatives	37
Figure 11. Prisoner's Dilemma Payoff Matrixes	38
Figure 12. Schematic Representation of Protection Motivation Theory (reproduced from Steg and Vlek 2009, p. 108)	50
Figure 13. Relationship between (Flood) Risk Characteristics (reproduced from Raaijmakers et al. 2008, p. 312)	51
Figure 14. Decision Problem: Objectivity and Subjectivity	55
Figure 15. Conceptual Model	57
Figure 16. Function Shapes	69
Figure 17. Risk Scores Correlation (risk seeking → risk averse)	76
Figure 18. Mean Total Investment in the Fund across the 6 Games, for each Experimental Session	79
Figure 19. Investment Behaviours and their Relative Frequency	80
Figure 20. Total Investment vs. Expected Payout	81
Figure 21. Mean Normalized Deviation from Optimum for different Order and Malignancy Treatments	84
Figure 22. Malignancy in Previous Round and Uncertainty (no significance estimates)	86

List of Tables

Table 1. Taxonomy of Game Theory (based on Finus 2001, p. 10).....	32
Table 2. Variable Formalization.....	63
Table 3. Lottery Characteristics	74
Table 4. Option Evaluation	74
Table 5. Risk Partial Scores	74
Table 6. Indifference Probabilities	75
Table 7. Descriptive Statistics: NormDevOpt	80
Table 8. Descriptive Statistics for each Uncertainty Treatment: NormDevOpt	82
Table 9. Descriptive Statistics for Different Measures of the Dependent Variable and Functions	86
Table 10. Questionnaire Results.....	87

List of Inserts

Box 1. Decision Theory: Analytical Distinctions.....	9
Box 2. Axioms and Theorems of Probability Theory	14
Box 3. Dimensions of perceived riskiness (based on Steg and Vlek 2009, p. 109).....	51
Box 4. Experiment Summary	67

Abbreviations

UN	-	United Nations
IPCC	-	Intergovernmental Panel on Climate Change
UNFCCC	-	United Nations Framework Convention on Climate Change
COP	-	Conference of the Parties
GSD	-	Global System Dynamics project
GHG	-	greenhouse gas
CPR	-	common pool resource
PG	-	public good
PD	-	Prisoner's dilemma

1. Introduction

1.1. Climate Change as an Observable Phenomenon

During the last decades, humanity has witnessed a gradual increase of the mean surface temperature of the Earth, known as “global warming”. Perhaps one of the most obvious effects of this phenomenon is the perturbation of the natural water cycle, which results in different precipitation patterns. The sum of modifications that are observed in the long-term statistics of weather and atmospheric conditions is called “climate change”. While climate change can also occur regionally, the term is usually used in contemporary environmental policy discussions to refer to the global changes that occur as a result of global warming.

The subject of, maybe, still too much political debate, climate change is one of the most urgent environmental concerns that humanity needs to address, as many of its consequences are irreversible and possibly detrimental to human life. A comprehensive overview of these impacts was given in 2007 by the Intergovernmental Panel on Climate Change (IPCC), in what is called the “Fourth Assessment Report” (IPCC 2007a). Since this report (abbreviated AR4) we know that “warming of the climate system is unequivocal”, with a total global mean temperature increase from 1850-1899 to 2001-2005 of 0.76°C [0.56°C to 0.92°C]¹ (IPCC 2007b, p. 5). Also, complementing the information provided by the IPCC, one of the latest scientific reports on climate change that came out as a result of a Scientific Congress held in Copenhagen in March 2009 shows that ecosystems and contemporary societies are highly vulnerable to increases in temperature beyond 2°C as compared to the pre-industrial era values (Richardson et al. 2009).

The good news, however, is that we also now know with a confidence of 90% that most of the observed global warming is attributable to the accumulation in the atmosphere of greenhouse gases (GHG) as a result of human activity since the mid-20th century (IPCC 2007b, p.10). This means that, in theory, we should be able to decrease the rate of warming and possibly reverse the trend until we once again reach the temperatures of the pre-industrial era. In practice, though, this is a much more complicated issue, because a significant reduction in greenhouse gas emissions would have broad consequences for current “Western” lifestyles and the economic aspirations of billions in presently low-income regions, as well as for the entire economic chain of production and consumption. Briefly put, this involves large short-term costs: for corporations these could range from high capital investments in clean technologies to cessation of production, while, for individuals, they could translate into increased taxation and product prices or, in some cases, quite radical changes in lifestyle and consumption patterns.

Current scientific knowledge advises us that a business-as-usual scenario in terms of GHG emissions and subsequent climate change could lead, in the long term, to a chain reaction of major social and environmental disruptions (Richardson et al.

¹ The range values for both the years and the temperature increase represent the uncertainty at 90% confidence level, i.e. there is an estimated 5% likelihood that the value could be above the range indicated and 5% likelihood that the value could be below that range. (see IPCC 2007b, p. 2, footnote 5)

2009). Hence, a precautionary approach is desirable. Nevertheless, it is difficult to estimate what the exact consequences will be and how future costs of inaction compare to the discounted present costs of action. In addition, due to the complex nature of the climate system and the large time scales associated with climate processes and feedbacks, projections about future developments are highly uncertain (Pittock 2009). By contrast to some other global environmental issues which have been successfully addressed at an international level (e.g. the ozone layer regime), it is these *uncertainties* that seem to make the climate change problem a really malign, difficult to solve, one. Giving up the comfort provided by preserving the status quo in order to avoid a risk of uncertain magnitude is not an idea that most people would readily embrace². And of course, this only adds to the difficulties already posed by the mere issue of *social coordination* in any type of collective problem.

1.2. Climate Change as a Policy Problem

Within the policy arena, people's interpretations of what climate change means vary significantly both across individuals and by comparison to the consensual scientific facts. Also, the information regarding the scientific uncertainty and the way in which this has been communicated to the large public has led people to evaluate what they know quite differently, resulting in a wide array of perceptions and views vis-à-vis the threat of climate change.

The very plurality of views supports the idea that a single interpretation and strategy are illusory (De Vries and Petersen 2009). As a consequence, there is no single, collectively-agreed upon, "optimal decision" that can solve climate change as a policy problem. Adaptation and incremental actions based on latest insights will probably be, for many organizations (governments, firms, NGOs) as well as individual citizens/consumers, the preferred way of dealing with this problem.

It might be reasonable to say that the broad collective objective is to prevent catastrophic climate change, but even with this in mind further debate arises on the optimal choices, depending on the meaning that individuals associate with this desiderate and on their understanding of the cause-effect chain. For some, solving the climate change issue is a minimal-cost strategy with lots of nuclear power and carbon storage; for others, it is an ethically motivated drastic life-style change. The prevailing view among UN-organizations and most governments is to make climate change policy part of the quest for "sustainable development", a notion that assumes, along the lines of the *ecological modernization discourse*, that environmental sustainability and economic development are compatible (Blowers 1998).

A consensus on practical steps to be followed is hard to reach, as also demonstrated by the weak results of the negotiations during the last UNFCCC Conferences of the Parties (COPs) and the failure to sign a global and comprehensive agreement on climate change. Yet in this vast ocean of opinions and interests, decisions are being made every day at an individual level, triggering collective consequences.

² Several psychological studies suggest that most people are ambiguity averse and prefer to avoid situations that are difficult to understand (see e.g. Pulford and Colman 2006, Pulford and Coman 2008)

When the private and public interests are in opposition we speak about a *social dilemma*. Within the context of climate change, distinguishing between *private* and *public* is problematic. In a very restrained form, the term “private” could be used to refer to those interests that concern solely an individual, while “public” would refer to what is common to more people. However, when it comes to UN negotiations, a “private” interest would concern a whole nation, by opposition with a global, “public” one. Ultimately, it can be said that many government officials participating in international negotiations do nothing more than represent the collective interests of the citizens that appointed them. When they judge a problem from a national perspective and with less consideration of the global issues, they oppose to a public interest their nation’s private concerns. Overcoming this social dilemma and focusing on the greater public good becomes an imperative for establishing a climate change regime.

Even within the optimistic paradigm of sustainable development, tackling climate change involves individual short-term compromises in order to achieve the long-term collective goal of stabilizing atmospheric CO₂ concentrations. At the country level, it is easy to monetize these compromises, but at an individual level, the currency for the associated economic costs might sometimes be a subjective measure of preference and utility rather than dollars or euro. A key challenge for governments becomes then to get an idea about how individual citizens perceive the issue of climate change, how they value the costs and benefits of individual and collective action, and which set of actions they consider feasible (cf. GSD project³).

To the extent to which citizens can influence high level politics (especially in democracies) and, from the other side, governments are interested in their opinion and policy support, understanding decisions made at the individual level might prove to be crucial in designing effective climate policies. While previous research has looked both into the individual vs. collective aspect (social dilemma theory), as well as the behaviour under risk and uncertainty one (decision theory), these two dimensions are seldom studied together and even less so within the context of climate change. It is exactly this gap that the present study attempts to fill in.

1.3. Research Objective and Main Question

The main purpose of this project is *to further clarify the policy problem of climate change, by exploring individual economic behaviour in a (dynamic) decision context that resembles that of climate change, with an emphasis on the challenges posed by uncertainty.*

Formally, our main research question is:

What roles, within the climate change context, do risk and uncertainty play in individual economic decision making and behaviour?

By answering this question we hoped to gain further insight into the main obstacles and opportunities for solving the climate change policy problem.

³ See <http://www.globalsystemdynamics.eu/> for details.

1.4. Research Framework

Because the policy problem of climate change is such a complex one, being ultimately dependent on individual decisions and strongly related to individual perceptions and actions, our attempt was to address the above mentioned question from an interdisciplinary perspective. Therefore, as it will become evident in the following chapters, our theoretical orientation combines knowledge from various fields, ranging from philosophy to psychology and experimental economics. Time and space limitations have of course constrained us to being very selective with the topics presented here and sometimes we had to dedicate more attention to certain disciplines than others. When this was the case, we gave priority to the theories or fields of research that have had a longer tradition in science. This was not to disregard the valuable contribution of newer theoretical advancements, but rather to provide a basis on which further insights could be added in future studies.

The research framework that we used in order to achieve our objective consisted of three main steps and is represented visually in Figure 1:

Step 1: Understanding Decision-Making

First, we tried to understand the basic mechanisms that underlie decision-making (and its behavioural outcomes) both in individual and collective situations. This was done through desk research, by reviewing the literature on decision and game theory. Since we found each of these two fields to be equally relevant to the problem of climate change, we dedicated one chapter to each of them. This research step corresponds to Part A in Figure 1 and the main findings are reported in Chapters 2 and 3.

Step 2: Understanding the Climate Change Decision Situation

Furthermore, we attempted to assess the current level of understanding of climate change as a policy problem and identify the particularities of this decision situation, as compared to other global environmental issues. Also, we looked for clues in the literature on individual behaviour in decision situations particularly framed as climate change. This was done by reviewing some relevant studies on social dilemmas and collective action as well as previous climate change related experiments (part B in Figure 1). Chapter 4 presents the main findings.

Step 3: Exploring Economic Behaviour in a Climate Change Decision Experiment

Next, based on parts A and B, we used a grounded theory approach in order to develop a simple behavioural model particularly suited to the decision context of climate change. In line with our research question, one of the most important variables we brought into discussion here was *uncertainty*. Having formulated a few hypotheses about the expected behaviour of people, we then explored some of the links in our model through an economic game/experiment⁴ in the lab (Part C in Figure 1). Thorough details about the model, the testing methodology and the results will be provided in Chapters 5 and 6. Our main findings are presented in Chapter 7.

Lastly, we started this journey with the ambition to use the outcomes of this experiment as basis for outlining a few recommendations for policy makers involved in the climate change issue. Of course, the realization of this secondary goal was to

⁴ The terms “game” and “experiment” are used here interchangeably.

depend greatly on the actual results. In Chapter 8 we reflect on our findings from the perspective of our research objective and discuss the extent to which policy recommendations can be made.

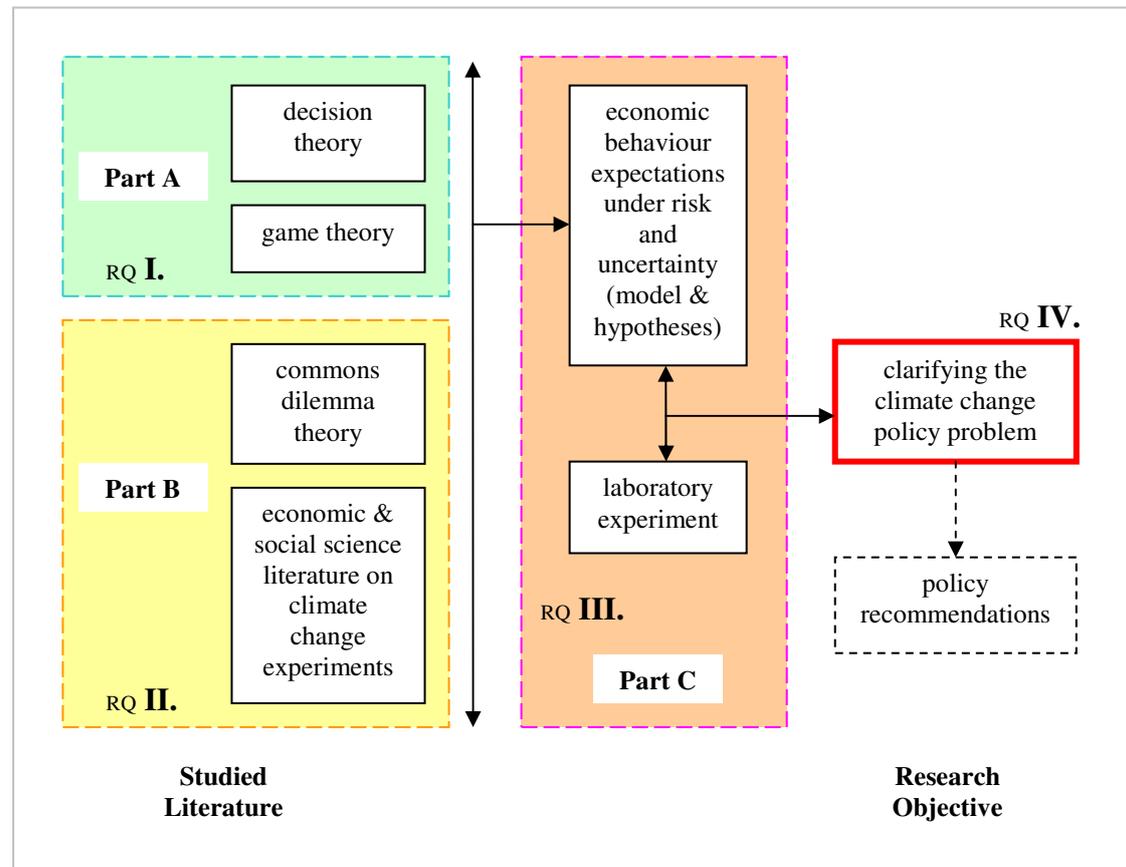


Figure 1. Research Framework

For each of the main steps we formulated some additional research sub-questions:

RQ I. What do we learn from decision and game theory about decision-making under risk and uncertainty?

RQ II. What do we learn from commons dilemma theory and previous climate change studies about human economic behaviour and how does this relate to the characteristics of climate change as a decision situation?

RQ III. Based on these insights, what are some useful hypotheses (i.e. which would add to our understanding) about human behaviour vis-à-vis climate change that could be derived and what is an adequate way of testing them?

RQ IV. Which hypotheses have explanatory power and what can we say about the climate change policy problem based on our results?

These questions have served as guidance for conducting our research, but also as the basis for structuring this report. Step by step, we will try to address each of these questions in the following pages and relate back to them when adequate.

2. Decision Theory

2.1. The Philosophy of Decision Making

Decisions are part of our every day lives. Whether we talk about a person who is wondering where to spend the next vacation or about a company evaluating whether or not to invest part of its capital in a new country, *making a decision* involves a *process* of choosing among different *options* or alternatives, as well as an *outcome* of this process, namely the *decision* or the *choice* made. At least this is the simplest depiction of decision-making that we might all be familiar with. On a closer look, however, the mere concept of *decision-making* already makes several philosophical assumptions: first, an ontological one that *alternatives* do exist, which collectively compose the *option space*; second, a cognitive one that the *decision agent* (be it an individual, a group or an institutional entity) perceives or is aware of such alternatives; third, a deterministic one that *the decision*, as reflected by a certain effect or disposition to act⁵, is the result of the decision-making process only; fourth, a formal theoretical one that the relationship between decision-making as a process and decision as a product of this process can be abstracted into some sort of *decision rule*.

To explicate these assumptions, consider a situation where a student, Paul, has to travel to school and decides to ride his bike. First of all, the fact that we even consider this a decision implies that, from the position of neutral observers, we agree that travelling from point A to point B is possible in multiple ways: by bike, by bus, by foot etc. Secondly, we assume that when considering his ways to travel to school Paul is aware of the existence of at least *some* of these options. This is not to say that the entire *option space* is immediately visible to him – he might, for instance, not even conceive roller-skating all the way to school as one possible alternative. Thirdly, when contemplating Paul's already made choice – i.e. to ride his bike – we assume that this is only as a result of a decision-making process that Paul has engaged in and not some external coincidence controlled by something we could call fate or predestination. In this simplification, no assumption is yet made on what exactly determined Paul to choose the bike over other alternatives, but an identity of *choice* and actor *intentionality* is indeed implied. Last, but not least, the fourth assumption refers to the fact that whatever caused Paul's choice might be formalized into a *decision rule*. He might have analyzed what was the fastest way to get to school given his specific circumstances, he might have tossed a coin, he might have relied on past experiences or he might have simply followed the advice in his daily horoscope – all these are viable rules.

Although we tried to keep these assumptions to a minimum, the discussion is already becoming complex when one considers such questions as: What are the limits of self? What constitutes intentionality? When can we say that a decision was made in full knowledge and acquiescence by the actor, and when that it was induced by a third party or made subconsciously? Since such issues belong to the realm of philosophy, no definite answers can be given.

⁵ We cautiously make this distinction here between “effect” and “disposition to act” to accommodate for what Weirich (2004) calls “mistakes in execution of plans”.

But even new interrogations arise as we pursue to establish causal relations between the decision process and its outcome. Having accepted that Paul's decision to take the bike is intentional, what is the reason that prompted him to choose this option over another? What is the decision rule that he applied and, again, was he aware of his application of the rule or not? Bringing into the discussion *rationality* or *preferences* further complicates the picture, because defining these concepts proves, in turn, to be problematic. Add to this the idea that the selection among different decision rules is itself a meta-decision, and we found ourselves in an overwhelming situation that resembles to the space between two mirrors that reflect each other infinitely!

In line with the philosophical tradition of scepticism perhaps everything is ultimately questionable, but this does not render futile all attempts at understanding the surrounding world. In Bertrand Russell's (1912, pp. 249-250) words, as uncertain as such questions and answers might be, "[they] enlarge our conception of what is possible, enrich our intellectual imagination, and diminish the dogmatic assurance which closes the mind against speculation".

2.2. Field Overview & Structure of the Chapter

Many are those who seem to have believed in this adage, because despite the above mentioned difficulties a vast body of literature on decision making currently exists. Scholars have tried to make sense of human behaviour in decision situations by proposing theories that would serve different objectives. Traditionally, one of the biggest questions to answer has been: "How can we make optimal decisions?", implying the existence not only of ideal actors in ideal circumstances, but also of a unique recipe that could lead everyone to similar (and ideal!) results with mathematical precision. As a critical response, newer theories have either tried to remove or loosen some of these idealizations or they adopted a more empirical approach by analyzing how decisions are made in practice.

There are different ways in which the many theoretical variations can be categorized, based on their philosophical grounds, purposes or basic assumptions. Classification is however controversial, as assigning one theory to one box is often a matter of interpretation. While it is not our purpose here to explain such analytical distinctions nor to map out the entire field of decision theory, an overview of some of the concepts encountered is given in Box 1. Nevertheless, it might be interesting to note that the term *decision theory*⁶ is sometimes used to refer strictly to a corner of this vast constellation of theories, namely to the normative and more mathematically-formal aspects (see Byron 1996, Weirich 2004).

It is exactly from this corner that we will begin our journey in the following sections, gradually expanding to more relaxed theoretical variations. After sketching the basic framework of a decision situation and explicating the key concepts we will present some of the main assumptions of what we call *traditional decision theory*. This includes *expected value* theory, but also the extensions to it brought about by economists with their *rational choice* and *utility* theories. Finally, after having considered the critique of the traditional approach, we will move on to reviewing

⁶ We will use this term to refer to the entire theoretical field, not only to some more formalized versions.

some of the newer theoretical developments that try to address the identified limitations.

Consequential vs. non-consequential decision theory

According to Weirich (2004) consequentialist approaches evaluate decisions in terms of possible outcomes or worlds – rationality consists in assessing and comparing the consequences of the different options and making a choice according to some optimization rule.

On the other hand, non-consequentialists place the emphasis in the discussion on the choice itself, instead of its outcomes. Some might consider choice to be an expression of rational values, while others could assess a decision based on the procedures that were employed in making it, rather than in the outcomes themselves*. In this case rationality could be seen, for instance, as the product of reason as opposed to that of passions.

As Seip and Wenstøp (2006) outline, consequentialist approaches tend to be anthropocentric – especially in environmental decision-making – and focus on results to the detriment of notions such as “right” or “wrong”. It is thus important to acknowledge that other kind of ethics exist that could be relevant in some situations.

Normative vs. descriptive/explanatory decision theory

The distinction between normative and descriptive decision theories is quite straightforward: “normative theory gives rules for what a *rational* decision maker *should* do in various situations [, while a] descriptive theory aims at describing how people *in fact* (rationally or not) make decisions in different situations” (Gärdenfors and Sahlin 1988, p. 5).

Descriptive theory is sometimes referred to as explanatory theory and stays at the heart of neoclassical economics, as it looks at the behaviour of people and assumes that agents are rational, in order to explain and make sense of their preferences. At the opposite side, while normative theory starts with the beliefs and preferences of the people in order to tell them how to act, it does not necessarily require that they will live up to these expectations (Morris and Ripstein 2001).

Traditional decision theory (see Section 2.4) is often interpreted as both normative and descriptive (Weirich 2004).

* for a list of consequentialist and non-consequentialist authors, see Weirich (2004).

Box 1. Decision Theory: Analytical Distinctions

2.3. An X-Ray of a Decision Situation

If we are to discuss decision theory it is necessary to understand first what the basic framework of a decision is. What does it mean to be in a *decision situation*? Leonard J. Savage, one of the earliest theorists of what we call the “traditional approach”, gives the following classic example:

“Your wife has just broken five good eggs into a bowl when you come in and volunteer to finish making the omelet. A sixth egg, which for some reason must either be used for the omelet or wasted altogether, lies unbroken beside the bowl. You must decide what to do with this unbroken egg. Perhaps it is not too great an oversimplification to say that you must decide among three acts only, namely, to break it into the bowl containing the other five, to break it into a saucer for inspection, or to throw it away without inspection. Depending on the state of the egg, each of these three acts will have some consequence of concern to you...” (Savage 1972, pp. 13-14).

From Savage's example we learn that three elements are essential for defining a decision situation:

- a set of *alternatives* (or *decision variables*), A , where the particular alternatives are denoted a_1, a_2 , etc.;
- a set of possible *states* of the world (also *state variables* or *states of nature*), S , containing s_1, s_2 , etc.;
- and a set of *outcomes* (or *consequences*),

where we denote a particular outcome as o_{ij} , if it is a result of choosing alternative a_i when the state of the world turns out to be s_j . These elements can be represented in a *decision matrix* like the one in Figure 2.

		States of the world			
		s_1	s_2	...	s_n
Alternatives	a_1	o_{11}	o_{12}	...	o_{1n}
	a_2	o_{21}	o_{22}	...	o_{2n}
	...				
	a_m	o_{m1}	o_{m2}	...	o_{mn}

Figure 2. Decision Matrix

All is simple and clear up to this point. The problem arises from the fact that an actor that finds himself in a decision situation: a) can only pick one alternative; b) doesn't know what the state of the world is; and c) since the outcome will have an impact on his life, he usually wants to make the "right" choice. Broadly speaking, decision theory, in all its versions, asks two general questions: 1) What is the "right" choice⁷ for a particular individual given a particular situation? and 2) What is the process behind an actor's choice (for descriptive versions of the theory) or what should it be (for normative versions)?

As Gärdenfors and Sahlin (1988, p. 1) rightly note, the common core of (rational) decision making theories consists of: 1) wants or desires that determine how we value the possible outcomes and 2) information or beliefs about how the world is and how our actions will influence it. But what are these wants, information and beliefs and what role do they play in decision-making?

If we think back about the example of a student who has to decide how to travel to school we can see where such concepts come into the picture. We propose the following simplified scenario: It is Monday morning and Paul has to be at school in half an hour, as he has to give a very important presentation. Because he is wearing

⁷ The "right" choice can be defined both in consequentialist and non-consequentialist approaches to decision-making: in the first case it would refer to the choice that produces the "best" consequences, in the latter it could be the choice that is "best" in itself.

his best suit he was planning to take the bus, but when he exits the house he realizes that he had lost his wallet and he has no money for the ticket. To make things even worse, he also doesn't have an umbrella and the sky is very cloudy and it seems like it is going to start raining soon. Paul has now two options: he can either take his bike while running the risk of getting to school wet if it rains, or he can borrow his parents' car, despite the fact that he really doesn't like to drive. The decision that Paul is now faced with involves his preferences and the information or beliefs that he has about the chances of rain during his trip.

In terms of how wants are related to information and beliefs, the assumption here is that if he knew it rained, Paul would feel compelled to drive as he cannot under any circumstances show up for this important presentation with his clothes soaked and this is more important for him than his discomfort when driving. Otherwise, he would just take the bike because he *prefers* this option over driving, independent of any other considerations. In both cases there would virtually⁸ be no decision to take at all. But the *uncertainty* over a state of the world (rain/no rain) which is *known* or *believed* to affect the outcome is what calls for a rationalization of the decision situation.

By "rationalization", we do not refer here to a process that necessarily uses the principles of rational choice theory (to be discussed in the following section), but rather to one that can be abstracted into a decision rule and thus motivate the choice, whatever that is. In the most general sense rationalization is simply an exercise of inductive logic that asks for internal consistency. As Weirich (2004) explains in his introductory chapter, there might be disagreements on the principles of rationality across different versions of decision theory, but not on this type of *instrumental* rationality in itself. This means that the baseline assumption is that, whatever the decision, there is always some way to justify it. The concept of *preference* is used to define "an ordering relation imposed on two objects or states of affairs" (Byron 1996, p. 17). By introducing this notion, decision theory focuses on the "hows" of the action, not on the "whys": it does not ask why Paul doesn't like to drive or why he can not show up for the presentation with his clothes wet, but only tries to prescribe, describe, or explain his choice based on these givens.

In order to elucidate this further, let's follow a similar example proposed by Carter and Price (2001) and try to sketch Paul's decision matrix for the situation described above (see Figure 3). Nothing in decision theory tells us whether Paul *should* prefer one alternative over the other, nor even if his choice is the cause or rather the effect of his preference (see discussion about consequentialism vs. nonconsequentialism in Box 1). Instead, the different variations of decision theory will diverge in their evaluation of rationality because of topic, scope or assumptions (Weirich 2004). Some might argue that it is not always possible to know the outcomes, let alone order them based on preferences, that the agent's goals or beliefs might change with the context, or that attitudes to risk and uncertainty might influence the decision process in unexpected ways.

⁸ By this we assume that Paul would immediately act in line with his most desired (certain) outcome, where: Paul's preference to look good for the presentation is greater than his preference to avoid driving. Thus if he knew for sure that it would rain, he would take the car, otherwise the bike.

		States of the world	
		<i>rain</i>	<i>no rain</i>
Alternatives	<i>bike</i>	Miserable - Can't give the presentation anymore - Guessed wrong	Happy & relieved - Guessed right - All's well that ends well
	<i>car</i>	Annoyed, but satisfied - Had to do something against his will - Was able to give his presentation	Depressed - Guessed wrong - Did something unpleasant for nothing

Figure 3. Paul's Decision Matrix

In the following sections we will present some of these accounts of rationality by referring back to Figure 3 where appropriate.

2.4. Traditional Decision Theory

2.4.1. Delineating the Field

As mentioned earlier, we use the term *traditional decision theory* to refer to that body of literature concerning human behaviour in decision situations that has originally been developed by mathematicians, statisticians, philosophers and, later on, economists. The common conceptual heart is derived from *probability theory* and is based on notions such as those of *expectation* and *expected value*. In addition to this core, the concepts of *subjective probability* or *expected utility* are often used to describe individual input in the decision-making process. Different definitions are given to these terms, but what is common to all the theoretical variations that we cluster within this category is the attempt to mathematically formalize the relationship between an outcome, its probability of occurrence, as well as the decision agent's preferences and attitude towards risk.

Traditional decision theory, sometimes often referred to as *formal decision theory*, can be interpreted both as a normative and as an explanatory (or descriptive) theory. In its normative aspect, it can be regarded as a branch of logic, developed from Aristotle's study of syllogisms (Weirich 2004, p. 4), where people are advised what the best act is, given their preferences and beliefs (Morris and Ripstein 2001). On the other hand, the explanatory interpretation is most predominant within neoclassical economics where the assumption is that people's preferences can be read and their economic behaviour modelled (Morris and Ripstein 2001). While it is true that the normative-explanatory distinction is not as clear cut as it might seem and that

the names given to the different branches of decision theory are often used inconsistently in the literature, it is often this explanatory interpretation that the term *rational choice theory* refers to.

As far as rationality is concerned, within traditional decision theory this is divorced from morality, in the sense that no claims are made about what the decision agent *should* desire. As Weirich (2004) well observes, this means that there is no distinction made between what is rational and what is reasonable or justified. Individual goals are private and given, which makes this theory a liberal one, accommodating different conceptions of “good” (Byron 1996). It is only the chosen means that can be analysed in terms of rationality, and *optimization* is here the golden rule: finding the most effective way of achieving one’s goals, based on one’s knowledge and beliefs (Morris and Ripstein 2001).

The questions that remain are how to weigh the effectiveness of an alternative when we do not know to what outcome it will lead, and how to rate the different possible outcomes considering that they always involve the risk of not being achieved. We will address these questions in the next two sub-sections by introducing some key concepts of decision theory.

2.4.2. Dealing with Uncertainty: Probabilities of Events

Probability theory, as a branch of mathematics, deals with some abstract concepts that attempt to measure uncertainty in problems of statistical inference and decision. In other words, it tries to give a numeric answer to a question such as: what are the chances that a particular (uncertain) outcome would occur?

The term *event* is used in probability theory to define “any set, or class, of possible outcomes” (Winkler 1972, p.5). Here the word “outcome” has a different meaning than the one we used when we looked at the elements that constitute a decision situation in the previous section. These are not the outcomes of a path chosen by a decision agent, given a certain state of the world, but the outcomes of a path chosen by nature (i.e. a state of the world). An event is an expression of what interests us with regard to the possible states of the world. In our example with Paul, we could be interested in the event of “rain” or “no rain” or “rain after 3 p.m.”. Of course, it is more difficult to mathematically formalize such statements about weather, than it is to discuss, for instance, the probabilities of different outcomes for a roll of a die – this is an issue we will get back to later on.

An event can be *elementary* – when it cannot be further decomposed into “smaller” events – or *compound* (Winkler 1972). For instance, when we roll a die, the event “occurrence of 1” is an elementary one, while the event “occurrence of a number that is larger than 3” is compound, because it can be decomposed into the events “occurrence of a 4”, “occurrence of a 5” and “occurrence of a 6”⁹. The set of all possible elementary events is called an *event space* or *sample space* (Winkler 1972).

Mathematically speaking, it is possible to define the probability of a certain elementary event, based on several axioms (conventions) presented in Box 2. Furthermore, a number of theorems can be deduced (see also Box 2) that allow to

⁹ Winkler (1972) also gives a similar example when he discusses the concept of “event space”.

perform operations with known probabilities in order to determine those that are unknown.

Axioms

Winkler (1972, p. 7) gives the following informal statement of probability axioms:

1. The probability of an event E , written $P(E)$, is nonnegative.
2. If S denotes the set of all possible events (the sample space), then the probability of S , $P(S)$, is equal to one.
3. If two events E_1 and E_2 are *mutually exclusive* (that is, they cannot *both* occur), then the probability that at least one of the two events will occur is the sum of the individual probabilities $P(E_1)$ and $P(E_2)$:

$$P(E_1 \cup E_2) = P(E_1) + P(E_2), \text{ which can be "translated" to}$$

$$P(E_1 \text{ or } E_2) = P(E_1) + P(E_2)$$

Theorems

Though not exhaustive, here are a few theorems that are useful when working with probabilities (Winkler 1972, pp. 7-10):

1. $0 \leq P(E) \leq 1$
2. If $E_1 \subset E_2$, then $P(E_1) \leq P(E_2)$
3. If *complement of E* denoted by \bar{E} is defined to be the event " E does not occur", then

$$P(\bar{E}) = 1 - P(E)$$

4. For two events that are *not necessarily mutually exclusive*,

$$P(E_1 \cup E_2) = P(E_1) + P(E_2) - P(E_1 \cap E_2), \text{ which is the mathematical notation of}$$

$$P(E_1 \text{ or } E_2 \text{ or both}) = P(E_1) + P(E_2) - P(\text{both } E_1 \text{ and } E_2).$$

This can be further extended to more events. For instance, with three events we get:

$$P(E_1 \cup E_2 \cup E_3) = P(E_1) + P(E_2) + P(E_3) - P(E_1 \cap E_2) - P(E_2 \cap E_3) - P(E_1 \cap E_3) + P(E_1 \cap E_2 \cap E_3).$$

An explanation of how these formulas are derived can be found in Winkler 1972, pp. 8-9.

Box 2. Axioms and Theorems of Probability Theory

The interesting question, especially from the perspective of decision theory, is what these numbers actually mean. Winkler (1972) mentions two main interpretations of probabilities: *the subjective interpretation* and *the frequency interpretation*. As Hacking (2006) shows, different other names are sometimes used to mark this same analytical distinction: subjective/objective, epistemic/aleatory, probability₁/probability₂¹⁰. Hacking (2006) himself classifies probabilities into *belief-type* and *frequency-type*. In our view, this latter dichotomy is preferable to the first one, not only because of symmetry reasons, but also because the terms "belief"

¹⁰ According to Hacking (2006), this distinction belongs to Rudolf Carnap.

and “frequency” already give an indication of what the difference between the two might be. For our purposes here, we choose however to keep the word “interpretation” because probabilities, as such, are of a single type, as defined by mathematics; it is only what we mean when we use them in daily life that is subject to divergence. In the following sections we will thus talk about the *belief* and the *frequency interpretations* of probability. We start with the frequency one because it is more related to the mathematical conceptualization, while the belief interpretation will be more useful when discussing other concepts from decision theory later on.

The Frequency Interpretation of Probability

According to this (classical) interpretation, the probability of an event is considered to represent the *relative frequency of occurrence* of that event in the long run (Winkler 1972).

When the assumption is made that all elementary events that comprise a sample space are equally likely, the probability of an event E can be determined by looking at its relative frequency in the space S :

$$P(E) = \frac{\text{number of elementary events comprising } E}{\text{total number of elementary events in } S} \quad (1)$$

If we consider the toss of a coin, for instance, we make the assumption that the occurrence of any side is equally likely. Then, based on the formula above, the probability of getting, say, “heads” is $\frac{1}{2}$, i.e. 0.5. As Winkler (1972, p. 11) points out, this is the relative frequency of occurrence of an event in a *sample space*, not “in the long run”. However, if an experiment is repeated many times, the theorem known as the *law of large numbers* states that the relative frequency of occurrence of a particular event will tend to be close to the probability of that event. Moreover, the relative frequency of occurrence of the event E , r/n , is more likely to approach the probability $P(E)$ as the number of repetitions n increases (Winkler 1972, p.12):

$$P\left[\left|\frac{r}{n} - P(E)\right| \geq \varepsilon\right] \rightarrow 0, \text{ as } n \rightarrow \infty, \quad (2)$$

where ε is any arbitrary small positive number, representing the margin of error. To put it more simply, the following theorem expresses the same idea: “As the number of trials increases, the accuracy probability¹¹ approaches 1. Relative frequencies tend to converge on probabilities.” (Hacking 2006, p. 197)

In other words, if we toss a coin four times, it is possible that we would obtain “heads” three times. This means that the relative frequency of occurrence of “heads” is 0.75, which is obviously different from the mathematical probability (calculated relative to the event space) of 0.5. According to the law of large numbers, the more we toss the coin the closer the relative frequency of obtaining “heads” will be to 0.5. This law is particularly relevant because it can also apply to real-life situations where the symmetry arguments are no longer applicable or plausible. Consider, for instance, a biased coin. If we toss it repeatedly for a large number of times and obtain “heads” in 70% of the cases, we may say that the probability of the event “heads” is 0.7. The mathematical 0.5 probability does not adequately describe reality anymore, as the

¹¹ Accuracy probability refers to the probability that r/n is within ε of $P(E)$.

assumption of equally likely occurrences is broken by the very physical properties of the coin. At the same time, a probability of 0.7 tells us that in the long run the relative frequency of “heads” will stabilize around the value of 0.7 (Hacking 2006).

Within the frequency interpretation of probability, an induction step is made where statistical regularity of historical observations is used to estimate the likelihood of an uncertain future event. In Winkler’s (1972, p. 13) words:

“The relative-frequency interpretation of probability, then, seems intuitively reasonable and empirically sound. If an uncertain situation is repeated a large number of times, the frequency of an event’s occurrence should somehow reflect the probability of the event. Most people think of this phenomenon as a «law of averages». It is important to note, however, that no matter how many repetitions of the uncertain situation are considered, the actual relative frequency of an event need not *exactly* equal the probability of the event. If a fair coin is tossed 1,000,000 times, the number of occurrences of heads is expected to be close to, but not necessarily equal to, 500,000.”

The same applies for the biased coin in our example: if we toss it for another 1,000,000 times, the relative frequency of “head” will not necessarily be 700,000. Nevertheless, if the number of past observations for which the relative frequency was calculated is high enough, it is reasonable to assume that in our new 1,000,000 trials the relative occurrence of “heads” will also converge to 70%, and thus it makes sense to put an equal sign between the relative frequency and the probability. Also, the larger the initial set of trials, the more precise our frequency estimation was and the smaller the deviation from 0.7 will be in a repeated experiment.

Given the above, relative frequencies are useful to determine probabilities when we have no other information about reality – for instance, when we doubt whether a coin is or isn’t fair. However, note that the law of large numbers concerns repeated independent trials with constant probability p for the event E , also known as *Bernoulli trials* (Hacking 2006). This could mean, for example, that the coin will not be changed in the middle of the experiment and, more generally, that the all other conditions will be identical for all trials. In some real-life situations it might happen that the exact conditions of a trial are difficult to reproduce or it might simply be impractical to have a long series of trials. That’s why, when trying to estimate the probability of a future event, other approaches might be more useful.

The Belief Interpretation of Probability

Practical experience teaches us that there are many events where the relative-frequency interpretation of probability is not at all useful for dealing with uncertainty. These events are usually unique and while it is possible to think of them in probabilistic terms, there are no underlying phenomena that can be subjected to a large number of trials in order to determine a certain statistical regularity. Statements such as “The odds are 1 in 100 that he will attend the meeting,” or “I will probably be

late for dinner,” use the notion of probability in a totally different way than the ones discussed so far. Consider the following statement taken from a BBC news report¹²:

“The scientists think that Venus may once have held copious amounts of water on its surface. But it is likely the solar wind removed most of it during the first billion years or so after the formation of the Solar System.”

In contrast to the previous two statements which were making predictions about the future, the idea of “likelihood” is used here to formulate a theory about a unique event in the past, but in all cases it is a certain degree of confidence that is expressed. This confidence can have many sources; hence, there are different theories of “belief-type” probabilities. For instance, *logical* or *interpersonal* probabilities describe the likelihood of a hypothesis relative to evidence (Hacking 2006, p. 142), like in the Venus example above. On the other hand, *personal* probabilities refer to the speaker’s degree of belief regarding the occurrence of an event, thus being expressions of individual judgments (Hacking 2006).

The inherent “subjectivity” of these probabilities might make such statements seem irrelevant for dealing with uncertainty under the constraint of rationality that traditional decision theory requires. When a fair coin is tossed (and it is certain that the coin is fair), with a probability of 0.5 of getting “heads”, one might argue that it doesn’t matter if an individual believes that the probability of obtaining “heads” in the next trial is, say, 0.9, because the “objective” probability is still 0.5. However, if we are to study decision making, it is exactly the belief interpretation of probability that we need to understand because it deals with the likelihood of a single trial, of a unique event. This does not mean, though, that there is no way to rigorously assess such subjective probabilities. Traditional decision theory, for instance, uses personal probabilities and, under the assumption of individual internal consistency, has developed some well defined tools for measuring them. This is something that will be discussed in more detail later on. For the moment it suffices to acknowledge that, as opposed to the frequency interpretation of probability, the belief interpretation can be applied to *any* uncertain situation, thus making it extremely useful in practice (Winkler 1972).

The Fine Line between the Two Interpretations: Information and Beliefs, Unique vs. Repeated Decisions

Not only do belief probabilities extend the range of situations where we can attribute a numerical value to the degree of uncertainty, but often they occur into decision making as expressions of individual confidence in an observed frequency based probability. For instance, a person might know from the weather forecast report that the probability of rain on a certain day is of 65%. It might be reasonable to assume that this probability had been calculated by meteorologists by observing a long series of similar atmospheric conditions to those of the day in cause, and the statistical relative occurrence of the event “rain” in those past situations was taken. This does not mean, however, that an individual in a decision situation where this information is needed would automatically use this probability as such. Instead, an entire set of

¹² Amos, Johnatan (2009). “Probe hints at past Venus ocean”. *BBC News – Science and Environment*, July 14, 2009. Retrieved online at: <http://news.bbc.co.uk/2/hi/science/nature/8149361.stm>. Accessed July 15, 2009.

experiences and motives might lead one to assign a different personal probability to the particular event. Perhaps one knows that the weather agency that made the forecast is not very reliable, and will assign to the same event a probability of 40%, just by looking at the sky. Thus, additional information beside the frequency probability might alter the way in which we use this so-called “objective” probability of 65% rain.

Another issue that frequency based probabilities raise concerns the inductive step that is needed when assuming that what is true for repeated decisions will also be true for unique ones. Consider the following simple bet situation: an individual is tossing a coin and he will win 1€ whenever the event “heads” occurs and he will lose 1€ whenever the event “tails” occurs. If he were to play this game for a large number of times, given the frequency interpretation of probability and the law of large numbers discussed before, both events will occur about half of the time, i.e. with probability of 0.5. For an infinite number of trials, the relative number of wins and losses would cancel each other out, so this bet would make no sense. This does not mean though that this game is necessarily a “no-gain-no-loss” one if played for a limited number of times. When played for, say 1,000,000 times, it is very possible that the winning side would show up 450,000, while the losing one 550,000 times, leaving the gambler in debt! A more optimistic person, however, could focus on the possibility of winning 50,000€, if the reverse was to happen.

In real-life we are confronted with single trials or, at most, with finite series of trials. If we think of a repetitive experiment, presumably under identical conditions, such as the toss of a coin, it might seem reasonable to use the observed relative frequency as an indication of probability over a finite series of trials, but this is in itself a “subjective” interpretation. According to the law of large numbers, relative frequency converges on probability when $n \rightarrow \infty$, so implying the same for any finite number of trials n is nothing more than a matter of belief. When we say that for *one particular toss* of coin the probability of getting “heads” is 0.5, because we have data confirming that the relative frequency of “heads” and “tails” was 0.5 and thus that coin is fair, we actually apply *the frequency principle*:

“If S is an individual event of type E , and the only information about whether E occurred on a trial of a certain kind, on a certain chance setup, is that on trials of that kind on that setup the frequency-type probability $\Pr(E)=p$, then the belief-type probability of S is also p . (Hacking 2006, p. 137).

The word “only” is here very important. Also, this principle holds only as long as we accept the assumptions that are needed for the frequency interpretation of probability (Winkler 1972), such as: independence of trials, identical conditions, definition of “long run” (“approaching”) etc. For the reasons discussed earlier, personal probability might be different from the (known) frequency one when:

- we find some of these assumptions unreasonable – e.g. contrary to the independence assumption, we think that the fact that we had 10 “heads” in a row indicates that a “tail” will soon occur;
- we have more/other information about the situation – e.g. we know (or we believe) that the historical records based on which the relative frequency is calculated have been intentionally falsified.

This is why, as Winkler (1972) acknowledges, the belief interpretation of probability can be regarded as an extension of the frequency interpretation.

Moreover, as we have already mentioned several times, there are also situations where the frequency probability is not known or cannot be calculated because the event is unique. In these cases, belief probabilities can also be used as a replacement. For instance, one might argue that we can never talk about frequency probabilities for the event “rain” because no two days are ever the same in terms of atmospheric conditions and thus the “identical trials” prerequisite for applying the law of large numbers is not met. In such a situation, the mathematician might recommend taking 0.5 as the probability of “rain”, by calculating it as the relative occurrence of this event *in the sample space*. This might be indeed a reasonable thing to do if no other information about the world and the causes of rain was known. But on a day when the sky is clear, in the absence of any other “objective” probability, it doesn’t make much sense to say that the probability of rain is 0.5, does it?

When we need to assign probabilities to real events, we will never use the mathematical abstraction (see formula 1.1.) as such, but the frequency or the belief interpretations of it. But in decision making situations it is a fine line that we’re walking and we might sometimes switch between the two interpretations with an ease that could be confusing at times. Leaving aside the problem of how to distinguish between the two, traditional decision theory assumes that it is at least possible to quantify the individual degree of confidence that results from this amalgam of information and beliefs.

Risk, Uncertainty, Complete Ignorance

So far we used the word *uncertainty* in its more general sense of “doubtfulness”, “incertitude”, to refer to the indetermination of the future, to describe the fact that one cannot know for sure what path nature will choose. We will now reduce this sense to a more specific one, which is used within decision making theory. Following from the previous discussion on probabilities, a distinction between risk and uncertainty is often made in the literature.

A decision is considered to be made under *risk* when “full” information about the world exists, i.e. the exact probabilities – under the frequency interpretation – are known (Gärdenfors and Sahlin 1988). For instance, if we want to extract a playing card from one complete deck we know enough about the composition of the deck to be able to say without any doubt that the probability of this card to be red is 0.5 (again under the frequency interpretation). On the other hand, a decision is made under *uncertainty* when no information is available whatsoever about the states of the world (Gärdenfors and Sahlin 1988).

In reality, though, it is difficult to say that there is full or no information about a certain issue; instead, many degrees of partial information exist (Gärdenfors and Sahlin 1988). Moreover, because under the belief interpretation of probability it is always possible to assign probabilities to certain events, some authors regard this analytical distinction as being artificial (see Winkler 1972). However, in newer versions of decision theory, to be discussed later on, it becomes clear why such a differentiation is worth preserving: people behave differently in situations where they have precise information about the frequency probabilities of an event as compared to the situations where they have to rely solely on their beliefs.

This is why perhaps a better way to distinguish between *risk* and *uncertainty* is by following the frequency/belief dichotomy made earlier. *Risk* would then describe a situation where “objective” probabilities can be determined based on historical data, while *uncertainty* would refer to events which are unique or for which limited information is available, compelling the use of belief probabilities instead (Carter and Price 2001). For risk situations the accuracy of the probability is quite high, while under uncertainty it is lower and an accuracy range could be specified (Carter and Price 2001).

Finally, *complete ignorance* defines a state in which there is no prior information about the probabilities of an event, yet the decision maker does not want to or cannot make a subjective estimation (Carter and Price 2001). Such a situation is also sometimes referred to in the literature as *strict uncertainty* and many decision rules are proposed, one of the most useful being Savage’s *minimax regret* (see Seip and Wenstøp 2006). This is similar to the famous wager of Pascal on the existence of God and states that one should always choose the alternative that brings the least possible regret. Another way to deal with such a situation is to regard it as a strategic interaction between the individual and nature. Under the assumption of complete ignorance, we can use tools from game theory to understand what the best choice is. This is something that we will come back to in the next chapter.

To conclude, the following quote summarizes well the difference between the three concepts discussed here:

“At one extreme, uncertainty is characterized by a known probability distribution; this is the domain of decision under risk. At the other extreme, decision makers are unable to quantify their uncertainty; this is the domain of decision under ignorance. Most decisions under uncertainty lie somewhere between these two extremes: People typically don’t know the exact probabilities associated with the relevant outcomes, but they have some vague notion about their likelihood.” (Tverski and Fox 2000, p. 116).

2.4.3. Dealing with Outcomes: Expectations, Payoffs, Utilities

In a decision situation different acts have different consequences also depending on the actual state of the world. As I was mentioning earlier, traditional decision theory assumes that a choice will be made by considering the information and beliefs about the state of the world, and with regard to the desired outcomes. But what are the desired outcomes and how to evaluate them relatively to the chance of achieving them?

The concept of *expectation* plays a central role here. According to the simplest definition, an *expectation* is a weighted average (Winterfeldt & Edwards 1986). Every act might involve some costs and benefits and an expectation refers to an attempt to somehow analyse these all together. This might appear as a sophisticated notion for decisions where the outcomes are represented by verbal descriptions, such as the ones in Paul’s decision matrix (Figure 3) discussed earlier. It is however very straightforward when we have numerical values as outcomes.

Simple Numerical Outcomes: Expected Value

Consider that an individual is offered the following bet: he may roll a die and will win 1€ if he gets “six face up” and will lose 50 cents otherwise. Two questions

arise here: 1) If he were to play this game over and over again, what would be the average payment for every trial? 2) Is this an advantageous bet, i.e. should he accept it or not?

The monetary values of the possible outcomes are called *payoffs* and can be represented in a *payoff matrix* (Figure 4) where net wins and losses are represented by positive and, respectively, negative numbers. Note that if the player doesn't accept the bet, the outcome is 0, regardless of the state of the world – he doesn't win nor lose anything.

		States of the world	
		“six face up” ($p_1 = 1/6$)	other ($p_2 = 5/6$)
Alternatives	accept bet	1€	-0.5€
	not accept bet	0	0

Figure 4. “Six Face Up” Bet - Payoff Matrix

In order to answer the first question, the payoffs have to be weighted by the probabilities of obtaining each: in this case $p_1 = 1/6$ for the event “six face up” and $p_2 = 5/6$ for all other possible events in the sample space (in this case, any of the other five faces of the die). This is called an *expected monetary value* or, more simply, an *expected value (EV)* and is calculated based on the following formula:

$$EV = \sum p_i v_i, \quad (3)$$

where p_i is the probability of every outcome i , and v_i is its corresponding value (payoff). In this case, the expected value of accepting the bet is:

$$EV_{bet} = \frac{1}{6} \times 1 + \frac{5}{6} \times (-0.5) = \frac{1}{6} - \frac{5}{12} = -\frac{3}{12} = -0.25$$

This means that if a person were to play this game repeatedly, on average he would lose 25 cents per trial. On the other hand, if the bet is not accepted, the expected value is:

$$EV_{nobet} = 0€$$

The second question is much more difficult to answer. In the long run, this is a losing bet, since EV_{bet} is negative. Given the assumption that getting nothing is better than losing something, one should decide not to accept this bet, by comparing the expected values of the two alternatives (EV_{bet} and EV_{nobet}). However, the expected value says something about repeated decisions, not unique ones. Thus, it can be easily argued that the calculated value of -0.25€ has no normative relevance for the decision making, because it has nothing to do with the outcome of a single trial – the latter can only be 1€ or -0.5€ (also see Winterfeldt and Edwards 1986). Accepting the bet or not becomes, in this case, a matter of risk attitude. In the following section we will

introduce a new concept that elegantly solves this problem, by taking into account the value that a particular outcome has for the decision-maker.

Complex Outcomes: Utilities

As Keynes stresses out, “in the long run we are all dead” (Keynes 1923 qtd. in Peterson 2008, p. 111). That is why, when dealing with unique decisions the agent’s preferences for certain outcomes and attitudes to risk are much more relevant than the expected value. The concept of *utility* is used as a theoretical construct to describe numerically such preferences. The basic idea is that people associate a certain value to every possible outcome that is not necessarily the same or proportional to the monetary payoffs (Carter and Price 2001). For instance, in the example above, one might value more the small possibility of winning 1€ than the certitude of not losing anything if the bet is refused, thus the *expected utility* of accepting the bet could be higher than that of not taking the bet.

Within the traditional interpretation of decision making, the basic assumption is that people have clear preferences regarding the possible outcomes. These preferences can be quantified as numerical values and may include “all that the agent cares about, not just self-interest” (Morris and Ripstein, 2001, p. 6). *Utility* becomes thus the currency for evaluating the outcomes, where it is understood as a non-linear function of money. In some version of decision theory not only are decision agents considered to be able to assign numerical values to outcomes independently of the risks involved, but they are also expected to be complete and consistent in their ordering of these outcomes by preference (Gärdenfors and Sahlin 1988; also further discussed in Section 2.4.4.).

If we accept the assumption that for each individual a utility function exists and can be elicited, then we can use it in ordering the preferred outcomes for each state of the world in the same way as we did with simple numerical values. Consequently, the concepts of *expected utility (EU)* and *subjective expected utility (SEU)* will be useful.

Expected utility was first introduced by von Neumann and Morgenstern (1944) and, like the expected value, can be easily calculated as:

$$EU = \sum p_i u(o_i), \tag{4}$$

where p_i is the probability that a certain outcome o_i would occur and $u(o_i)$ is the utility of that outcome o_i .

The probabilities p_i that are used to calculate the expected utility of a decision situation are, in this case, assumed as known and “objective” (Camerer 1995).

The concept of *subjective expected utility*, is a more refined version of EU, first developed by Savage (1972[1954]). Its main merit is that it extends its application, by using subjective probabilities derived from people’s preferences. Under SEU theory, all “unknown” probabilities can be inferred from choice (Camerer 1995). Of course, the same formula as for EU applies (see Equation 4).

As we were saying earlier, *utility* has long been a very popular concept in economics, as it seems to elegantly solve many problems concerning how to deal with outcomes. It is quite intuitive to say that for someone not losing 10 € could be more useful than having a small chance to win 1000 €. However, if we refer back to Paul’s decision matrix in Figure 3, we will immediately see that there are situations when it is still difficult to order utilities: is it more “useful” to be miserable or depressed? In

the following section we will give an account of one of the most popular versions of decision theory that deals explicitly with methods to elicit utilities and then make the “rational” choice. We will also see what the underlying assumptions are, in what ways they are limited and why it was necessary to move on to different theoretic approaches.

2.4.4. Rational Choice Theory

We were saying earlier that *rational choice theory* is often interpreted as an explanatory branch of traditional decision theory because it reveals the people’s preferences. At the same time, it is also in a way prescriptive, as it uses these preferences, as well as the subjective probabilities (where the case), in order to define what is “rational” and most efficient.

Having emerged from economics, rational choice theory (also called Bayesian theory) uses sophisticated strategies¹³ to scientifically measure people’s preferences and utilities. Without insisting here too much on the methodology, we should retain that – within rational choice theory – neither the subjective probabilities, nor the utilities associated with each outcome are simply spontaneous beliefs that people might have. Instead, building the utility function and determining the probability estimates are done by asking the decision agents to state their preferences over uncertain prospects (see Peterson 2008). Once this step is completed, the subjective expected utilities are calculated for each alternative and the alternative that leads to the best outcome is selected as the “rational” choice.

Basic Assumptions

The main assumptions that rational choice theory makes, and from where all further debates arise, are nicely summarized by Gärdenfors and Sahlin (1988, pp. 3-5) as following:

A1 (Values of outcomes): The values of the outcomes in a decision situation are determined by a utility measure which assigns numerical values to the outcomes. [...]

A2 (Values of alternatives): When determining the value of a decision alternative, the only information about the decision maker’s wants and desires that is exploited is the utilities of the possible outcomes of the alternative. [...]

A3 (Information about states): A decision maker’s beliefs about the states of the world in a given situation can be represented by a unique probability measure defined over the states. [...]

A4 (Probabilistic independence): For all states and all alternatives, the probability of the state is independent of the act chosen.”

Furthermore, when he introduces subjective expected utilities, Savage adds an additional axiom, known as the *sure-thing principle*, which can be formulated as:

¹³ See Peterson 2008, Chapter 2, for a more detailed explanation of how this is done.

“The choice between two alternatives must be unaffected by the value of outcomes corresponding to states for which both alternatives have the same payoff.” (Gärdenfors and Sahlin 1988, p. 8)

As far as the methodologies employed in order to elicit utilities and subjective probabilities are concerned, these are also subject to certain implicit conditions. As Morris and Ripstein (2001) summarize, it is assumed that an agent can compare all alternatives (“completeness requirement”), that he will rank one ahead of the others in any given choice (“transitivity requirement”) and that all these choices do not contradict each other (“consistency requirement”).

All these assumptions are useful in the mathematical demonstrations that are needed to build a unitary theory of rational choice. Depending on what approach they are trying to justify, economists will provide their own interpretations of the given assumptions or only rely on a subset of them. Nevertheless, we should say once more that all variations of rational choice theory use *the principle of maximizing expected utility (MEU)* as a decision rule: “Among the acts you can perform at a time, perform the one whose utility is at least as great as any other’s utility” (Weirich 2004, p.20).

Limitations

It will not be our purpose here to make an extensive critique of rational choice theory, but rather to briefly present some of its limitations, as a justification for the more recent developments in decision theory.

In general, many discussions arise in the literature on the validity and interpretations of the assumptions presented above. If we look at A1, for instance, this states that it is always possible to determine not only the order of the preferred outcomes, but also how many times more an individual prefers one alternative over another. In our example with Paul, if we determined that “happy” gave him a utility of “4”, while “depressed” corresponded to a utility of “1”, we are not only suggesting that he prefers to be happy to being depressed, but also that this is exactly 4 times stronger. It is straightforward to see why such an assumption could not be easily accepted. Another point, regarding A3 this time, is that because the probability estimates are subjective, an actor could take any decision and still be rational. The arguments for refuting one assumption or another range from the abstract and philosophical to the mere practical ones, such as the famous Ellsberg’s paradox (see discussion in Gärdenfors and Sahlin 1988, pp. 12-13).

Within a normative interpretation, an important limitation of rational choice theory is its incapacity of being of actual help to decision actors confronted with uncertainty. When we deal with known values of the outcomes and probabilities, it is easy to apply the maximization rule, but when such values are less evident, it is often impractical to elicit utilities. Some critics of Bayesian approach decry exactly the fact that the preferences of the agents are the source of the recommended strategy, seeing this as circularity. In Peterson’s words (2008, p. 26):

“Bayesians put the cart before the horse from the point of view of the deliberating decision maker. An agent who is able to state preferences over a set of uncertain prospects already knows what to do. Therefore, a Bayesian agent does not get any new, action-guiding information from the theory.”

On the other hand, under a descriptive interpretation, this theory fails to account for all observable behaviour. People are often inconsistent in their choices and what might seem to them as rational under some conditions might seem irrational under other conditions (e.g. at a different point in time). Moreover, it takes a lot to assume that people will mathematically ponder every choice they make, even if this were really easy.

It is mostly for these reasons that newer developments of decision theory have tried to relax some of the idealizations of decision actors and rationality and have been more preoccupied to describing decision behaviour and finding possible patterns.

2.5. Alternative Theories

Having introduced most of the useful concepts in the previous subchapter, here we will only give a brief account of some of the later developments in decision theory.

As it has also been the case in the past, newer developments to the traditional model have been advanced by economic science. Behavioural decision theory, presented in Section 2.5.1, is perhaps the “mainstream” approach to decision making research these days, as it combines the methods and formalizations of experimental economics with social sciences. Because it is built empirically, behavioural decision theory needs no further philosophical grounding. Despite the fact that we live in a time where the constructivist paradigm is the dominant one, perhaps more normative non-consequentialist approaches should not be disregarded, as some of them bring into discussion interesting ethical and value-driven considerations. Section 2.5.2 is dedicated to one such theory that we thought might be relevant in the context of climate change.

2.5.1. Behavioural Decision Theory: New Developments

Most theories that developed as alternatives to the rational choice versions started from the idea that people have limited computational abilities and use simple “heuristics” when faced with decisions. This was first proposed by Simon (1955) in his influential paper “A Behavioural Model of Rational Choice”, where the *principle to satisfice* is suggested as a more realistic decision rule than optimization or MEU: the decision agent sets an aspiration level and adopts the first discovered option that meets that requirement. While it also starts from the individual’s desires, this principle is supposed to be closer to real behaviour, as it liberates the decision agent from making complex evaluations of all alternatives and possible outcomes.

In line with this proposition, psychological research has shown that in many cases the assumptions of individual rationality made by traditional decision theory do not hold. Economists initially discounted this evidence invoking methodological or implication arguments, but when they tried to replicate psychology studies by use of the methods of experimental economics, they found that indeed some results could not be explained by the traditional theory (Camerer 1995). Consequently, these empirical findings have gradually been included in the classical economic models.

The accumulated evidence on the systematic violations of EU and SEU theories or axioms has led to the development of what is called *behavioural decision theory*. From this continuously growing body of knowledge, we mention here two empirical findings that have become theories in themselves.

Prospect Theory

According to this theory developed experimentally by Kahneman and Tversky (2000) since 1979¹⁴, people make choices with respect to a *reference point*, as it is change and relative values that matter to them more than absolute values. This has two main implications:

First, the utility function is not only concave (and behaviour only risk averse), but rather changes its shape around that reference point (“reflection” effect), so that people will be risk-avoiding in gains, but risk-seeking in losses (see figure in Kahneman and Tversky 2000, p. 34) In other words, people are willing to take considerable risk to avoid losses, but not for extra gains over an already satisfying level. Of course, the reference point differs among individuals and cases.

Second, objective probabilities over the state of the world are weighted differently (transformed into subjective probabilities) when they are small from when they are large. In line with previous empirical research, prospect theory suggests that people will tend to overweigh low probabilities and underweigh high ones, with an inflection point between .1 and .3 (see figure in Camerer 1995, p. 621).

Ambiguity Aversion Theory

The concept of “ambiguity” has been introduced by Ellsberg (qtd. in Camerer 1995) to describe those decision situations where relevant information that could be known is not known. The classical decision problem proposed by Ellsberg involved an urn that contained 30 red balls and 60 black or yellow balls in unknown proportions out of which the agent has to draw one at random. Although not rigorously defined by Ellsberg himself, our understanding is that the concept of *ambiguity* overlaps with that of *uncertainty* as we defined it in a previous subchapter; it covers, thus, the grey area between risk and complete ignorance, where decisions have to be made based on subjective probabilities.

Ambiguity aversion theory, first tested empirically by Becker and Brownson (1964), states that people prefer a known-risk situation to an ambiguous one. More specifically, when confronted with different bets in the above-mentioned decision situation, individuals will violate the principles of rational choice theory in such a way that they avoid the ambiguity.

Formally, there were several ways suggested as of how to include this finding into traditional theory, but none found much interest in terms of empirical testing (see Camerer 1995 for details).

Rather than trying to fit ambiguity aversion into the traditional theory, current research in this field has focused mostly on further understanding this effect under different conditions. For instance, Liu and Colman (2009) studied choices between risky and ambiguous bets in single and multiply repeated choices. Their finding was that under the repeated choice treatment more decision makers selected the ambiguous gamble. In another study, Chakravarty and Roy (2008) separate between the effects of risk and ambiguity on choice behaviour in order to see, among other things, whether the “reflection” effect in risk (as studied by Kahneman and Tversky

¹⁴ 1979 is the year when this study was originally published; the 2003 reference sends to a volume in which the original paper has been republished.

2000) also extends to ambiguity. One of their general results is that people are ambiguity neutral over gains and mildly ambiguity seeking over losses (Chakravarty and Roy 2008).

2.5.2. Gauthier & Rationality of Commitments

So far in this account of decision theory we have insisted on what were the dominant views in the literature. Considering the historical and philosophical contexts under which they have developed, perhaps it is not by accident that the most important theories were also those that have emerged from economic science. The kind of rationality that traditional decision theory proposes is very well aligned with the positivistic and individualistic paradigm that has made its way into our society from the Enlightenment and until today.

Although we are about to jump on a very different path now, we insist on presenting here an interesting alternative view where the mere concept of “rationality” is put in a different light. While there are certainly other similar theories, we chose Gauthier’s point of view as a representative example of a totally different mindset, where the long-divorced rationality and morality are brought one step closer. We do this exercise because we believe that it is within this kind of approaches that we might find the key to many sustainability problems, including climate change.

The main argument of Gauthier is that morality puts a rational constraint on the pursuit of individual interests (Yi 1992). Traditional decision theory is focused on future outcomes, where preferences cannot be inferred from past actions. As Morris and Ripstein (2001) explain, commitments, on the other hand, are backward-looking by nature and keeping them could be, in fact, rational at times. In order to prove his point, Gauthier still uses the maximization principle inherent to traditional theories, but he embeds this in what he calls “dispositions to perform” (see Yi 1992). The philosophical implications and discussions from here onwards can become quite complex, so we will not go into details here and now. The main point is, however, that traditional decision theory is limited because it does not accommodate for the fact that prior agreed rules (thus “dispositions to perform”) could lead to a greater good than the direct and independent pursuit of one’s own interest, thus making it rational to act in conformity with these rules.

The significance of such an approach to rationality will become even clearer in a later chapter when we will discuss decision situations – such as the famous Prisoner’s Dilemma – that require cooperation between multiple actors.

2.6. Conclusions

After this long journey in the vast field of decision theory, the task is now to decide what we want to take along further in our study.

Perhaps one of the most important, yet trivial conclusions is that decisions are a matter of individual judgment and that there is no single rule that is applied systematically by all people in the same way. Even if we accepted a definition of rationality as traditionally imposed by economists, by no means could we say that this kind of rationality is indeed the norm that is followed when it comes to actual behaviour. While the traditional interpretation of decision theory could be useful – and even valued in its prescriptive form – in some situations, we should rather seek explanations of human behaviour where such theories are most at home: in the realm of social sciences.

Concepts such as risk and uncertainty are not easy to handle, to begin with, yet even less when not even consistency of thought and action can be expected from individuals. Later on, in our experiment, but also in real life climate change related decisions, we should expect people to be highly adaptive to the situation at stake. Whether their reactions will be psychologically biased in certain circumstances (cf. prospect theory), or will be based on ethical concepts, this is another question.

3. Game Theory

3.1. Game Theory as Interactionist Decision Making

Generally speaking, game theory studies decision situations where more individuals interact, thus promising to be applicable to any kind of social phenomena. As it is the case with decision theory, game theory can also be interpreted in a normative or descriptive way.

Traditionally, game theory developed as a branch of mathematics, first introduced by Von Neumann and Morgenstern in 1944 (Camerer 2004), that proposes analytical solutions to what *rational* players should do in a strategic situation. Whether the strategies recommended by game theory are regarded by theorists as a norm or simply as an exercise of logic, it is a matter of their view on rationality.

However, when it comes to the *prediction* ability of the theory, it has been empirically observed that people will often deviate in unexpected ways from the analytically determined “best” path. Hence, it is the mission of the newer field of *behavioural game theory* to systematically collect such deviations.

Based on this difference of approach between traditional game theory and newer versions of it, this chapter will follow a structure similar to that of Chapter 2, although we will not go into the philosophical accounts of rationality again.

3.2. Traditional Game Theory

In its traditional form, game theory preserves some of the axioms of decision making, but no probability estimations are made about the occurrence of different outcomes, which are all considered to be equally likely. The choice of a strategy is then dependent on other player’s possible strategies.

Having emerged from mathematics, game theory proposes some problems – called *toy games* – and tries to solve them analytically. These are usually storylines that describe a certain decision dilemma that might call for some strategic thinking. Often in the literature much debate arises on whether these toy games are a good representation of human interactions and whether people do or should behave as game theory tells. In our view, before even considering this kind of arguments, a distinction should be maintained between game theory as an analytical exercise and its applications and relevance to real-life. The toy games should be regarded merely as conceptual constructs, similar to chess or logic problems, where a well-established theory might in some cases propose an analytical solution. Whether that solution is (not) or should (not) be employed by decision-makers, it is a different story, which does not invalidate, however, the theory that recommended it. As Camerer (2003, p. 5) very nicely explains on a similar note:

“If people don’t play the way theory says, their behavior has not proved the mathematics wrong, any more than finding that cashiers sometimes give the wrong change disproves arithmetic”.

This confusion where an analytical theory has become to be understood by some as an attempt to model life comes from the fact that game theory had been long used by economists to prescribe what “rational” actors should do. How this happened and what are the implications, we will discuss later. For the moment, we would like to

confine game theory to its simple mathematical form and introduce some of the key concepts it deals with.

3.2.1. Key Concepts

In order to introduce some of the key concepts of game theory, let's go back to our friend Paul and take one classical toy game as an example: the *coordination game*.

Suppose Paul has to decide whether he wants to go to school by bike or by car, but this time the problem he is confronted with is not whether it will rain or not, but what will his friend, Ana, do. Paul knows that Ana has to make the same choice between riding the bike and driving. He would prefer to ride his bike when she is also riding her bike (he is sure that they would meet on the way) and to drive otherwise. Most of all, he would be unhappy if he were riding the bike alone. The same can be said about Ana's preferences. How can we analyze this situation using game theory and what should each of them do when confronted with this decision?

The Representation of Games

First, it would be useful to represent the story above in a simpler form. A problem like ours is usually summarized in a *payoff matrix*¹⁵, or what is called the *normal form of a game*. This includes the following elements (Gibbons qtd. in Finus 2001, p. 22):

- the players in the game;
- the strategy combinations, and
- the payoffs received by each player for each possible combination.

Following the lines presented in the decision theory chapter, payoffs could very well refer to monetary values, or utilities. A *strategy* in this situation would correspond to an *action*, or the choice that each player has to make. Note, however, that there is a distinction to be made between these two terms: if this game were to be played repeatedly, an *action* would refer to one instance of this decision situation, while a *strategy* could refer to a set of choices in the long run (Finus 2001).

Going back to our example, if we are to assign some payoffs to the problem above, the payoff matrix would look like the one in Figure 5:

		Ana	
		<i>bike</i>	<i>car</i>
Paul	<i>bike</i>	2, 2	0, 0
	<i>car</i>	0, 0	1, 1

Figure 5. Coordination Game: Payoff Matrix

When represented in normal form, the assumption is that the players in that game have to take their decisions (act) at the same time and independently, i.e. Ana and Paul are each at home and cannot communicate to each other. However, in some games players move in sequence and in those cases it is the *extensive form of a game*

¹⁵ Key concepts will appear in bold italics throughout this section.

that is more useful. This representation uses a tree-diagram to indicate which player moved first, such as the one in Figure 6. More formally, the extensive form of a game should include the following (Gibbons qtd. in Finus 2001, p. 44):

- the players in the game;
- when each player has to move;
- the actions available to each player at each decision node;
- the information each player has at his turn;
- the payoffs to each player at each possible end node.

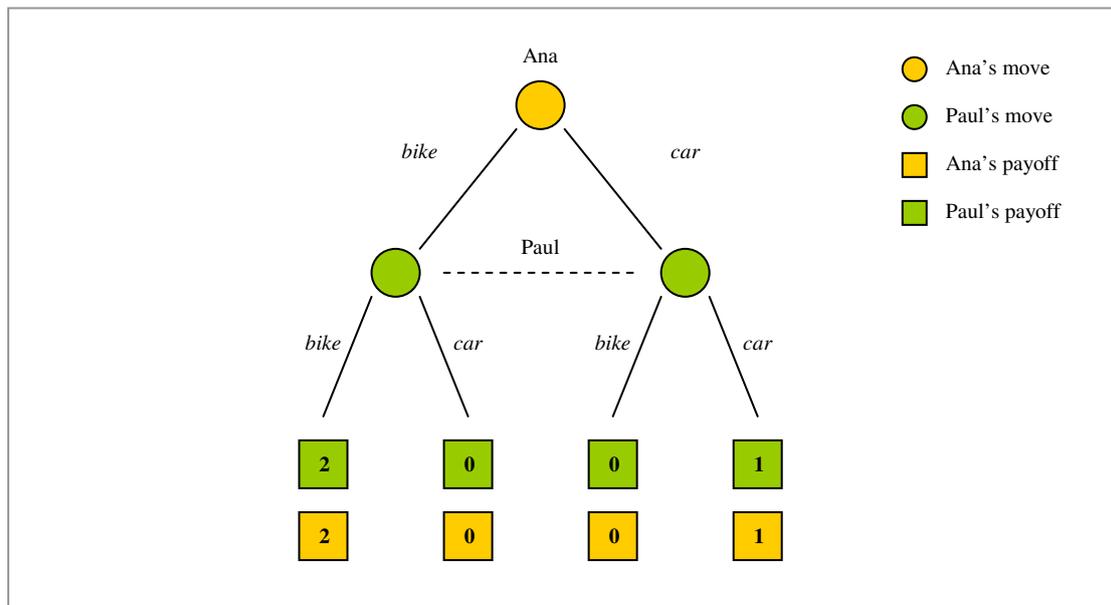


Figure 6. Cooperation Game in Extensive Form

The information element is symbolized through a dashed line. To preserve the conditions of the coordination game described before, we here connected Paul's two decision nodes. This means that, at the moment of his choice, Paul does not know what Ana had played at the previous move. Had he known this, then the dashed line would have been absent from the tree-diagram (Hargreaves Heap and Varoufakis 1995).

Taxonomy of Games

While in its original version game theory focused on very simple toy games, later developments to the theory have proposed increasingly complicated problems. Hence, contemporary game theory studies many different types of games. Before we proceed to discussing possible theoretical solutions, it makes thus sense to briefly mention here some of these categories. Table 1 presents a list of criteria and concepts as summarized by Finus (1997 qtd. in Finus 2001) to which we added some short explanations¹⁶.

The cooperation problem mentioned previously is an example of a game that would be characterized as following: it is *non-cooperative* as Ana and Paul don't have a previous agreement on what to do and even if they had, any of them could still freely break it; it is *non-constant sum*, because the players do not have diametrically

¹⁶ Our explanations are also based on Finus 2001, pp. 10-17.

opposed interests; it is a game with 2 *players*, that choose across a *discrete strategy space* (to take the car or the bike), with *complete information* and *simultaneous moves*. Since nothing is said about Paul and Ana having this decision problem every day, this can also be considered a *static game*. The other criteria do not apply here.

Table 1. Taxonomy of Game Theory (based on Finus 2001, p. 10)

1. Character of the game	(a) cooperative	Presumably, binding agreements can be enforced (e.g. coalition formation games).
	(b) non-cooperative	Agreements cannot be enforced, players pursue their own interests.
2. Cost-benefit structure	(a) constant sum	A player can only win at the expense of others (who necessarily lose).
	(b) non-constant sum	Cooperation can generate additional welfare; “win-win” situations might exist.
3. Number of players	(a) 2	2 players: solutions are easier to determine / cope with.
	(b) N	More players: generalization, mathematically more complex.
4. Strategy space	(a) discrete	The possible choices are discrete (e.g. choice to pick up the phone or not).
	(b) continuous	The decision space is continuous (e.g. choice of a tax level, this can be any number).
5. Time horizon	(a) static	One-shot game; the game is played only once.
	(b) dynamic: (i) finite, (ii) infinite	Repeated game. A dynamic game is finite when it is known with certainty when the game ends; it is considered infinite if it lasts forever or if the end of the game is not known with certainty.
6. Time dimension	(a) discrete	Time is counted in discrete intervals (e.g. a game has more <i>rounds</i> or <i>periods</i> and actions can only be taken at certain moments).
	(b) continuous	Actions can be taken at any time.
7. Time structure	(a) independent	Also called <i>repeated games</i> : a “basic” game is played repeatedly over several rounds. This does not also imply <i>strategic</i> time independence, as learning effects might occur from one round to another.
	(b) dependent	Also called <i>differential games</i> : the payoff at any point in time depends on the payoffs and actions taken in previous rounds.
8. Information requirement	(a) complete	All information is known to all players.
	(b) incomplete	At least one piece of information is not known to at least one other player.
	Generally, it is assumed that all players know the information allocation (i.e. if the game is of complete or incomplete information)	
9. Sequence of moves	(a) simultaneous	All players make their moves at the same time.
	(b) sequential	Players make their moves one after another.

Dominance, Equilibrium & Basic Assumptions

So far we learnt to characterize the toy game that we proposed, but we still didn’t show what game theory calls a *solution* to this game. In other words, under a normative interpretation, what course of action – if any – would game theory recommend?

In order to be able to answer this question, we need to use concepts such as that of *dominance* or *equilibrium*. Broadly speaking, a strategy is considered to be *dominant* “if it is a best strategy [...] regardless of the opposition’s choice of strategy” (Hargreaves Heap and Varoufakis 1995, p. 44). At the same time, an

a outcome is an *equilibrium* if “it is brought about by strategies that agents have a good reason to follow” (Hargreaves Heap and Varoufakis 1995, p. 45). But what is a “best strategy” and what is “a good reason to follow”? From here onwards, the discussion becomes complicated, because specifying what the equilibrium point of a game is depends largely on the assumptions made (Finus 2001).

Traditional game theory makes use of the following basic *assumptions* (see Hargreaves Heap and Varoufakis 1995 for a critical examination and further details):

- individual action is instrumentally rational (utility maximization); this basically means that the “best strategy” is the one that maximizes utility;
- there is a *common knowledge of rationality* (CKR) – each person is instrumentally rational and knows the same about their opponent;
- there is a *consistent alignment of beliefs* (CAB) – “no instrumentally rational person can expect another similarly rational person who has the same information to develop different thought processes” (Hargreaves Heap and Varoufakis 1995, p. 25);
- individual action is taken within the rules of the game.

At a conceptual level, the common knowledge of rationality is considered to be of *first-order* if each instrumentally rational player believes that the others are also rational in the same way. In other words, when Paul analyzes what his best choice would be, he assumes that Ana would follow a similar way of reasoning as he does. If Paul would change his strategy because he thinks Ana thinks what he thinks about what Ana thinks...etc. to *n* such steps, he would be using *n*th-order CKR (Hargreaves Heap and Varoufakis 1995). This kind of high-order CKR is useful in some special cases, as we will show later.

Going back to the notion of *dominance* and our cooperation game, if we examine the payoff matrix in Figure 4 we see that if Ana takes the bike, Paul would be better off biking; on the other hand, if she drives, he would also be better off going by car. The notion of *best reply* to other’s choice is used to indicate the alternative that gives the highest payoff in that case. It is customary to symbolize the best replies by using a circle for the row player, Paul, and a square for the column one, Ana, as shown in Figure 7 (Binmore 2007). In this case, no strategy is dominant because there is no one alternative for Paul that would maximize his utility, irregardless of what Ana does. Since the payoff matrix is symmetric in this case, the same can be said about Ana’s options.

		Ana	
		<i>bike</i>	<i>car</i>
Paul	<i>bike</i>	(2)	0
	<i>car</i>	0	(1)

Figure 7. Payoff Matrix with Best Replies for Paul and Ana

As far as the concept of *equilibrium* is concerned, we will refer here to two main types of equilibrium:

First, an *equilibrium in dominant strategy* is a combination of strategies that are both dominant for each of the players. This happens when each player can determine what his best strategy independent from what the other will do (Finus 2001). The cooperation game does not have such an equilibrium because there are no dominant strategies for neither of the players. The Prisoner's Dilemma (PD) game that we will discuss later is an example of a game with such an equilibrium.

Secondly, the concept of *Nash equilibrium* makes use of the CAB assumption and it refers to a combination of strategies that are all the best responses to the strategies of other players. In other words, two strategies are in a Nash equilibrium if no player has an incentive to deviate from his strategy provided that the other does not deviate (Finus 2001). For instance, in the cooperation game the pairs of strategies (*bike, bike*) and (*car, car*) are such equilibria because, if played, they confirm the expectations that the players have about each other.

It should be noted that in the case of a cooperation game, there is no unique solution to the problem of which pair of strategy will or should be played. Without further assumptions, traditional game theory cannot recommend a strategy to follow. One way out would be for Ana and Paul to communicate and coordinate on one of the two Nash equilibria. An important consideration here is also that, when given the chance to talk, out of the possible equilibria Ana and Paul should coordinate on the (*bike, bike*) one, as it renders more utility for both of them, i.e. it is *Pareto-efficient*. Another theoretical solution would be to assume 2nd order CKR and say that both Ana and Paul will always choose to ride the bike, precisely because they both think that the other will do the same in an attempt to achieve the Pareto-efficient equilibrium. Without such further assumptions, however, this game's solution is indeterminate.

Introducing Probability Distributions: Mixed Strategies vs. Pure Strategies

So far the ideas presented here were based on *pure strategies*. A pure strategy refers to a precise choice on how to play the game. In a one shot game, such as the coordination game, a pure strategy is confounded with a move: taking the bike and taking the car are both pure strategies. On the other hand, a *mixed strategy* involves playing a pure strategy with a certain probability (Finus 2001; Hargreaves Heap and Varoufakis 1995). To be more precise, a player plays a mixed strategy when he plays a pure strategy at random. According to Binmore (2007), mixed strategies were introduced in game theory as an extension to pure strategies in an attempt to ensure that all finite games have Nash equilibria.

In our coordination game example no solution could be found in *pure strategies*. Ultimately Paul and Ana would have to choose at random whether to take the bike or the car, with no guarantee that the other chose the same. The question that is asked is: how can one choose at random and still be rational? A *mixed strategy* implies that a player decides to make the choice between his pure strategies with a certain probability. A mixed strategy is, thus, a probability distribution $p_i = (p_{S1}, p_{S2}, \dots, p_{Sn})$, where $p_{S1}, p_{S2}, \dots, p_{Sn}$ are the probabilities assigned to each pure strategy, so that their sum is 1 (Finus 2001). An important observation here is that p_i is not the probability of a particular outcome, not the probability that a certain strategy will be played, but simply a probability distribution that player i takes in his mixed strategy.

The fact that a player assigns probabilities to his pure strategies, still does not answer the question of what the “best” distribution p_i is; how should a player distribute these probabilities so that he maximizes his utility? The concept of *Nash equilibrium in mixed strategies* is useful here, defined as a particular probability distribution where each player’s mixed strategy is a best response to the other player’s mixed strategy (Finus 2001).

We might recall from the previous chapter that a probability can be interpreted as a frequency. To illustrate this discussion about probability distributions and the concept of Nash equilibrium in mixed strategies, it is easier for our purpose here to imagine that Ana and Paul would have to play the same coordination game every day. Also, to make things simpler, we modify the payoff matrix so that it looks like the one in Figure 8, where the two players are indifferent between the two possible equilibria in pure strategies. If Paul had no idea about Ana’s choice, his best reply as a mixed strategy would be to take the car and the bike with equal probability, i.e. half of the times, because in the long-run he will be right in 50% of the cases. Thus, Paul’s probability distribution would be $p_p = (0.5, 0.5)$. At the same time, Ana would reason in the same way and her “best” bet would also be to drive half of the time and ride her bike the other half, where $p_a = (0.5, 0.5)$. Thus the Nash equilibrium in mixed strategies of this game is when Ana and Paul choose the mixed strategies of distribution p_p and p_a respectively.

In our example, determining the “optimal” probability distributions is quite straightforward, but more complex payoff matrixes require more advanced math that we will not present in this study. The main point here is that game theory has developed the tools to solve even apparently insolvable games. Conceptually, this might prove interesting, but it does not mean that mixed strategies are necessarily useful in real-life. For example, one important critical point that could be raised in this regard concerns the use of mixed strategies in one-shot games (also see Finus 2001).

		Ana	
		<i>bike</i>	<i>car</i>
Paul	<i>bike</i>	<div style="border: 1px solid black; display: inline-block; padding: 2px;">1</div>	0
	<i>car</i>	0	<div style="border: 1px solid black; display: inline-block; padding: 2px;">1</div>

Figure 8. Modified Payoff Matrix

3.2.2. Game Theory and Decision Making under Complete Ignorance

Now that we presented some of the key concepts in game theory, we would like to address an issue raised in the previous chapter and explain in this section how game theory tools can be for decision making under the assumption of complete ignorance.

First of all, decision making under complete ignorance is regarded as a game with two players: one player (the decision maker) has to choose an alternative, while player two (“nature”) “selects” a state of the world (Carter and Price 2001).

Let’s remember the example from the previous chapter where Paul’s dilemma was whether it would rain or not and what means of transportation he should consequently choose. For our purposes here we assume that somehow he was able to assign utilities to the different possible outcomes (happy and relieved, miserable, depressed etc. from Figure 3), so that these now constitute the *payoff matrix* shown in Figure 9:

		Nature’s “alternatives” States of the world	
		<i>rain</i>	<i>no rain</i>
Paul’s alternatives	<i>bike</i>	2	4
	<i>car</i>	2	1

Figure 9. Paul’s Decision: Payoff Matrix

Under the complete ignorance assumption, Paul will assign no probability to the events “rain” and “no rain” whatsoever. At the same time, “nature” is considered in this case to be indifferent between the alternatives (as it has no will or preferences on its own), hence its payoff would be 0 for any of the four outcomes. In this situation, the concept of *dominance* from game theory could recommend us a “best” course of action, where different alternatives are compared relatively to each other and the dominated ones are eliminated. Formally,

“an alternative a_k is said to dominate an alternative a_j if, for every possible state $[s_i]$ ¹⁷ alternative a_k is at least as good as alternative a_j ” (Carter and Price 2001, p. 305).

A quick look at the payoff matrix shows that, for Paul, the utility of taking the bike when it rains (i.e. of being “miserable”) is the same as the utility of driving when it rains (i.e. of being “annoyed but satisfied”). On the other hand, if it doesn’t rain, taking the bike will still lead to more utility (i.e. the utility of being “happy and relieved”) than driving (i.e. the utility of being “depressed”). Hence, the alternative of driving is dominated because Paul will always be better off if he takes the bike, regardless of the state of the world.

In this example, the payoff matrix was of such a kind that the *dominance* principle could nicely solve our problem. However, it might have also happened that Paul’s utility of taking the bike when it rains was 1. In that case, no straightforward solution could have been found.

Apart from the *dominance* rule, there are also other principles that could ground one’s choices, depending on context, the decision-maker and the payoff matrix. For an exemplification of these other rules, consider the matrix shown in Figure 10 and refer to our summary below (based on Carter and Price 2001, pp.305-308):

¹⁷ For consistency reasons, we used our own notation here and replaced the original θ_i by s_i .

		States of the world	
		s_1	s_2
Alternatives	a_1	1	6
	a_2	0	3
	a_3	4	2

Figure 10. Hypothetical Payoff Matrix with Three Alternatives

- **Maximin strategy** (“eternal pessimist”): consider what the worst possible outcome is for each alternative and then choose the alternative that has the maximum value of this minimum;
 - e.g. the worst outcomes are: 1, 0, and 2; then, a_3 should be chosen because it is the one that guarantees at least 2, indifferent from the state of the world.
- **Maximax strategy** (“eternal optimist”): for every alternative consider the best possible outcome and then choose the one that gives the maximum value of this maximum;
 - e.g. for the three alternatives, the best outcomes are: 6, 3 and 4; then a_1 should be chosen because it gives the biggest possible payoff.
- **Laplace principle** (principle of insufficient reason): assume that every state is equally likely and calculate the “expected payoff” for each alternative; then choose the alternative with the highest expected payoff;
 - e.g. if every state is equally likely, this means a probability of 0.5 for s_1 and 0.5 for s_2 . Then the expected payoffs are: $EP(a_1) = 1 \times 0.5 + 6 \times 0.5 = 3.5$; $EP(a_2) = 0 + 3 \times 0.5 = 1.5$; and $EP(a_3) = 4 \times 0.5 + 2 \times 0.5 = 3$. The chosen alternative should consequently be a_1 .
- **Hurwicz principle** (“degree of optimism”): the decision maker’s degree of optimism is defined as a number α , where $0 \leq \alpha \leq 1$ and for each alternative the *Hurwicz measure* is calculated as: $h_i = \alpha \{ \max_j p_{ij} \} + (1 - \alpha) \{ \min_j p_{ij} \}$, where p_{ij} is the payoff associated with alternative a_i and state s_j ; then, select the alternative with the highest value of h_i ;
 - e.g. suppose $\alpha = 0.7$; then $h_1 = 0.7 \times 6 + 0.3 \times 1 = 4.5$; $h_2 = 0.7 \times 3 + 0.3 \times 0 = 2.1$ and $h_3 = 0.7 \times 4 + 0.3 \times 2 = 3.4$, thus recommending a_1 as the best strategy.
- **Savage Minimax Regret**: for every possible payoff p_{ij} calculate the regret $r_{ij} = p_j^* - p_{ij}$, where p_j^* is the best outcome that could occur under state s_j ; then apply minimax on this matrix (find the maximum regret for each alternative and select the alternative that minimizes the regret).
 - e.g. under s_1 the regret column would be 3, 4, 0 and under s_2 it would be 0, 3, 4; next, the maximum regret for each alternative would be: 3, 4 and 4 respectively; finally minimizing the regret recommends a_1 as the best strategy.

While useful in practice, these other rules in fact stretch the borders of traditional game theory and the corresponding strategies are not as robust as the ones related to the concept of dominance. As Carter and Price (2001) rightfully note, when selecting one of these strategies the decision-maker is implicitly assigning probabilities to the outcomes, and thus breaking the assumption of complete ignorance.

3.2.3. Prisoner's Dilemma Game

In this section we would like to pay attention to a toy game that is particularly relevant for environmental problems. The so-called *prisoners' dilemma* (PD) refers to a game where two players each have to choose between a strategy to "cooperate" and another one to "defect"¹⁸. The original story refers to two prisoners who cannot be convicted unless one of them confesses and to whom the judge offers separately the following deal:

"If you confess and your accomplice fails to confess, then you go free. If you fail to confess but your accomplice confesses, then you will be convicted and sentenced to the maximum term in jail. If you both confess, then you will both be convicted, but the maximum sentence will not be imposed. If neither confesses, you will both be framed on a minor tax evasion charge for which a conviction is certain." (Binmore 2007, p. 5)

Theoretically, the payoff matrix of a PD game satisfies the $a > b > c > d$ condition, as depicted in Figure 11a. In Figure 11b we also give a numerical example of the payoffs associated with the above storyline. Note that the values are not utilities, but years to be spent in prison, thus the $a > b > c > d$ relationship is assumed to be inversed.

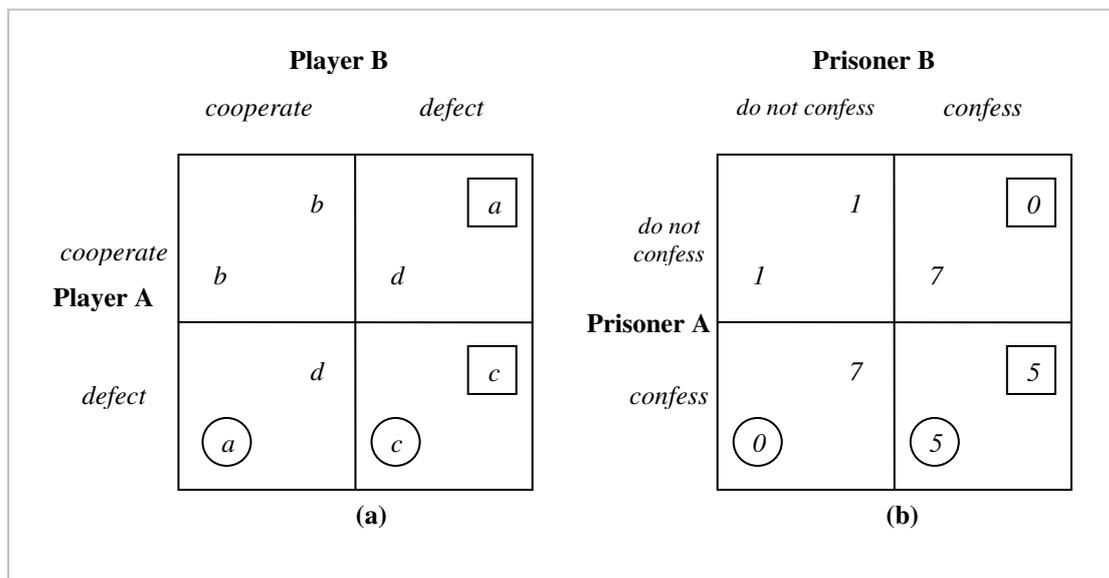


Figure 11. Prisoner's Dilemma Payoff Matrixes.

If we analyze the options of Prisoner A when Prisoner B does not confess, his best reply is to confess, as he will then be set free. At the same time, when Prisoner B does confess, A's best reply strategy is still to confess, so that he only spends 5 years in prison instead of 7. Because of the structure of the payoff matrix, the "confess" strategy is dominant for prisoner A. Since the matrix is symmetrical, the same

¹⁸ The strategies are called "dove" and, respectively, "hawk".

reasoning also applies for Prisoner B. The strategy pair (*confess, confess*) is then a Nash equilibrium in dominant strategies and a solution to this game.

In line with the basic assumptions of traditional game theory, rational choice compels both prisoners to confess. However, the dilemma of such a situation is that both prisoners would have been better off if they both remained silent as the strategy pair (*do not confess, do not confess*) is Pareto-superior to the equilibrium point.

While theorists saw in this dilemma a model of many real-life problems, the PD game only became a real subject of debate when economists tried to also put it forward as an explanation or prescription of actual human behaviour. In the normative sense, rational choice theory would require people to always defect in a PD game. However, game theorists cannot stress enough that the solution to the PD game should only be seen as a type of mathematical exercise, with no prescriptive or explanatory intentions:

“[I]t [is] just plain wrong to claim that the Prisoners’ Dilemma embodies the essence of the problem of human cooperation. On the contrary, it represents a situation in which the dice are as loaded against the emergence of cooperation as they could possibly be. If the great game of life played by the human species were the Prisoners’ Dilemma, we wouldn’t have evolved as social animals!” (Binmore 2007, p. 9)

In fact, empirically it has indeed been found that there are many PD like situations in which people are able to coordinate so that they achieve the Pareto-superior outcome (see e.g. Cadernás and Ostrom 2004). Moreover, even theoretically it is possible to find more efficient solutions if the game is, for instance, played repeatedly. The prisoner’s dilemma remains, nevertheless, an interesting way of formalizing and structuring real-life interactive decisions and a topic difficult to avoid when discussing environmental cooperation.

3.2.4. Advanced Game Theory

In terms of mathematics, we so far kept the discussion to a minimum level, by focusing on simple one-shot games and the most basic concepts associated with them. It is beyond our scope here to insist on more advanced game theory and to present theorems and demonstrations about equilibriums in more complex games. A comprehensive treatment of such games with application to international environmental cooperation is given by Finus 2001.

For our purposes, it suffices to briefly enumerate some general topics that more advanced game theory touches upon. First, one of the most straightforward extensions to simple toy games is to see what happens when they are played repeatedly. Mixed-strategies and sub-game equilibria are concepts that are then brought into discussion. Second, the time dimension is studied – different strategies arise when games are played for a finite or theoretically infinite number of times (the end is not known), and also when the time space is discrete or continuous. Finally, other directions of study include learning effects, the enlargement of the strategy space through issue linkages, and coalition formation (cooperative game theory).

3.3. Behavioural Game Theory

Following the same kind of developments as decision theory, traditional game theory raised questions regarding its relevance as an application to real-life strategic

interaction. Because formal game theory fails to predict or accurately describe observed human behaviour, a whole body of empirical tests were started in an attempt to determine possible patterns in deviations and to incorporate them in the traditional model. These formalized collections of observed behaviours constitute what is known as *behavioural game theory*. The idea underlying this pursuit is that a descriptive theory could actually help people faced with strategic choices, by telling them what they can expect in terms of other's behaviour. Ultimately, the goal is not to “«disprove» game theory (a common reaction of psychologists and sociologists), but to *improve* it by establishing regularity, which inspires new theory.” (Camerer 2003, pp. 20-21)

In the following subsection we would like to compile a list of some main empirical observations of behavioural game theory. This list will not be exhaustive though, but will only prepare the ground for later analysis. Some of the ideas presented here will be further developed in Chapter 4 of this study, where we will specifically discuss climate-change related games and experiments.

3.3.1. Methodological Considerations

Before we proceed to describing some of the main empirical findings of behavioural game theory, we should consider here a few methodological issues. Through its very nature, behavioural game theory involves decision-making and social interaction; hence, the object of study is a quite unpredictable and complex system. Also, most of the research in this field is experimental and so the question arises of how to ensure that the observations confirm indeed the existence of some regularity and are not just simple coincidences.

As Camerer (2003) notes, behavioural game theory comes up with different ways of testing the *joint* hypothesis that players have social preferences that value their own payoff *and* that they play game-theoretically. When – as it often happens – this joint hypothesis is rejected, the next challenge is to determine which of the parts of the joint hypothesis caused the rejection: was it that players had other considerations than their own good, or did they just not play rationally? How to ensure that this step is still possible?

Both of the concerns mentioned above can be dissipated, according to Camerer (2003), through the construction of careful and robust experimental designs. This implies, among other things, making sure that the subjects understand the rules and that they are highly motivated when performing the given tasks. To this, we would add that the exigencies imposed to other experimental fields of social sciences should also be applied to behavioural game theory.

3.3.2. A Few Empirical Findings

We now continue with a few interesting findings that might be particularly relevant within the context of our research and we group them in some thematic categories, following some suggestions by Camerer (2004).

First, there is enough evidence to construct a *theory of social utility* about how much of their own interest will players sacrifice for the sake of others. Within this category many experiments have shown that, when given a choice, people will tend to seek fair or equal outcomes rather than those that maximize their own profit (see e.g. Ambec and Ehlers 2008). However, the judgments of fairness will depend on context, culture and the decision's maker “*inequality-aversion*” (see e.g. Henrich qtd. in

Camerer 2004, p. 376). Other important findings here include the “*reciprocity theories*” that explain an observed correlation between player’s degree of cooperation and his/her expectations of cooperation by others (see Camerer 2004). An interesting extension here is the theory of “*rule-bound behaviour*”. As demonstrated by some studies, players will sacrifice themselves in order to punish or reward others’ behaviour, but when given the choice, they will also avoid finding out if their sacrifice is needed and just assume that it’s not (see Rabin qtd. in Camerer 2004, p. 380).

A second category of experiments concerns a so-called *theory of iterated reasoning and first-period play*. Iterated reasoning refers to the previously mentioned n^{th} order CKR and, in order to find the equilibrium point, much iteration is necessary for some games. However, experimentally it is observed that most people will stop after 0-2 steps of iterated reasoning in the first period of play (see Nagel qtd. in Camerer 2004, p. 380).

Third, a vast body of literature has led to the development of a *theory of learning*. Learning rules were studied in the lab, an interesting one being the “*experience-weighted attraction*”. This is a hybrid of reinforcement and belief learning models and it states that for each strategy the player associates a level of attraction and he then updates it throughout the game so that it reflects experiences (Camerer 2004).

Finally, there is also a research area that tries to determine if some of the behavioural theories of individual decision-making also apply for games. *Ambiguity aversion*, for instance, seems to still be present in games, where it manifests as a reluctance to take action when important information is missing (Camerer 2004).

Of course, the accumulated experimental data on choices in strategic interactions covers many more issues that we managed to briefly cover here. However, the above-mentioned theories should give a good impression of the kind of research that is done in this field. In the following chapter we will reconsider some of these findings and elaborate on those that are particularly applicable to our climate change decision context.

3.4. Conclusions

Having enlarged our view of decision theory with additional insights from game theory, it is now time to reflect on the utility of our undertaking within the context of this project.

As indicated in our research framework, the purpose of this part was to find out what the main lessons from decision and game theory are regarding decision-making under risk and uncertainty.

First of all, we learnt that if we are to build a theory of decision-making that has any explanatory power, this needs to be based on empirical findings and to incorporate psychology and social theories of utility. Behavioural decision theory and behavioural game theory are such attempts to build a more realistic, less idealized, picture of human behaviour.

Specifically in the context of climate change and the associated high level negotiations, traditional game theory has often been regarded as a useful analytical tool for understanding cooperation or the lack thereof in international environmental agreements (Finus 2001). Indeed, it is perhaps easier to depersonalize governments and devoid them of any other considerations than their own interest. The assumption

that governments will maximize their own welfare and act according to the rational choice postulate seems a bit more realistic when it is applied to an institution. Some might, thus, argue that it would have made sense to dedicate more attention to concepts such as those mentioned in the “Advanced Game Theory” section.

However, for the purposes of our study, we decided to disregard aspects like inter-country negotiation and strategies, to instead idealize about well functioning democratic systems, and focus on individuals as the ultimate source of decision making at the international level. This has to do with our own belief that societal pressure can constitute a source of change and high-level decision, but also with a choice that was already made explicit in our research question, i.e. to follow the lines of experimental economics and study individual economic behaviour.

Within this limited understanding of the climate change policy problem, we can retain from our literature review and use in later chapters the following few ideas:

- people will behave differently in situations of risk than in situations of uncertainty;
- people will take more risks in order to avoid a loss than to pursue a gain;
- situations of one-time decisions are likely to lead to different behaviours than those where the same kind of decision has to be made over and over again;
- learning effects occur from one decision situation to the other;
- if we are to understand and explain human behaviour in decision situations, we should also use insights from social sciences and remember that cognition or rationality cannot be separated from individuals’ mental maps, emotions, values and needs.

In the next chapters we will further specify some of these lessons within the context of climate change, gradually preparing to build our own conceptual model for addressing our main research question.

4. Environmental Commons Dilemma and Climate Change

In this chapter, we move on to our next sub-question that aims at specifying decision-making within the climate change context. As we have already established so far, in order to make sense of what drives human behaviour in the face of risk and uncertainty, we will need to look more in depth at social science literature. First, we bring into discussion the concept of *environmental commons dilemma*, which can be well applied to the climate change policy problem, and review some of the literature surrounding it. Next, we discuss the special characteristics of climate change as a decision situation/policy problem, as a step forward towards designing an economic experiment. Lastly, we summarize some of the findings of previous experiments and studies related to climate change, regardless of whether they particularly address economic behaviour or explore the relationships between perceptions, attitudes and behaviour.

4.1. Introducing the Concepts

A conceptual framework that is useful for understanding the environmental problems that we face today is that of *the resource/environmental commons dilemma*, which owns its popularity to an article written by Hardin in 1968. The *commons* are resources that are shared among a number of individuals, confronting each of them with the decision of how much to take from that resource or how much to contribute to its maintenance.

Generally speaking, the dilemma itself consists in a conflict between the “optimal” choices prescribed by the individual and, respectively, the collective rationalities. While on the individual level “optimal” can very much be a subjective matter, from the collective perspective the word is used here along the lines of the definition of sustainability given by the Brundtland Commission, so that it refers to a decision that “meets the needs of the present without compromising the ability of future generations to meet their own needs” (Brundtland 1987).

In a larger sense, a commons dilemma does not reduce itself to a simple choice between individual and collective interests, but has a more complex structure than that. From the perspective of the decision maker a commons dilemma raises several questions that have to be addressed simultaneously. Vlek and Keren (1992, pp. 250-251) distinguish between four types of trade-offs¹⁹:

1. A *benefit-risk* dilemma: if security, comfort and wealth are primary human desires, then to what extent should the level of achievable benefits be restricted in order to keep associated risks acceptably low?
2. A *temporal survival* dilemma: if the principal task of individuals, groups and organizations is to survive ‘now’, to what extent should present security, comfort and wealth be restrained in order to safeguard future survival conditions (which one may not live to see)? What amount

¹⁹ Adaptation of the original text.

of resources could be best devoted to short-term survival and how much should be allocated to long-term survival?

3. A *spatial survival* dilemma: if it is our principal task to survive ‘here’, to what extent should our local security, comfort and wealth be limited so as to secure more global survival conditions such as the quality of the seas or of certain forest areas?

4. A *social survival* dilemma: if one’s principal task is to survive in person (or as a group or organization), then to what extent should individual security, comfort and wealth be restricted in order to maintain collective survival conditions such as public utilities, education, transport and health care?

From an analytical point of view, decision and game theory provide us with a framework that can be useful in formalizing these four types of questions. As shown in Chapters 2 and 3, people are not to be expected to be aware of or necessarily ponder all these questions at once. However, the four categories proposed by Vlek and Keren (1992) provide a structure to the kind of considerations that a decision-maker might have in mind when confronted with a choice regarding an environmental commons.

At the other end of the individual-resource relationship, depending on the characteristics of the resource, there are two main types of dilemmas: *common pool resource dilemmas* and *public good dilemmas*. **Common pool resources** (CPRs) are characterized by the fact that it is difficult to exclude or limit people from using them – i.e. they are non-excludable – but at the same time one person’s use reduces the availability of the resource for others – i.e. they are rival in consumption (Ostrom et al. 1994). Examples include fisheries, forests or pastures. Similarly, **public goods** (PGs) are also non-excludable, but consumption by one individual does not affect the availability of the good to others. Public television and the atmosphere are traditionally regarded as examples of public goods (Ostrom et al. 1994). A tax for education or for roads is also a good example of a public good.

One way of analysing the commons dilemmas is by using game theory. Within this discipline a distinction can be made between *cooperative* and *non-cooperative* situations/games (see Section 3.2.1, Table 1): a game is in the cooperative category if engagements made by the players cannot be revoked (Finus 2001). Consequently, as Faysse (2005) points out in a review article, one dominant paradigm in the research of the commons is that of non-cooperative theory, as it is assumed that irrevocable rules are impossible or too costly, or that users are not able to create the conditions for such rules. Cooperative theory, on the other hand, assumes that people *are* able to establish rules and will do so as long as the benefits resulting from this shift are higher than the costs of design, monitoring and sanctioning (Ostrom 1999 qtd. in Faysse 2005). It is along the lines of the cooperative branch of game theory that most behavioural studies dealing with the commons have emerged.

In the following sections we will first present a game theoretic model that has been often used to describe the commons dilemma, and then we will discuss some empirical observations and explanatory models of human cooperation. But before we follow this path, we will first identify some of the special characteristics of climate change as an environmental commons dilemma.

4.2. Climate Change Particularities

Within the context of climate change and rising global temperatures due to greenhouse gas emissions, the use of the atmosphere as a sink seems to be a special case of a commons. On one hand, it resembles a CPR problem, where it is difficult to exclude actors from emitting harmful gases (i.e. use “clean air” units), yet the cost of this action is shared by everyone (i.e. “clean air” is rival in consumption). On the other hand, the structure of the CPR dilemma can be changed by “creating” more “clean air” units (e.g. by planting trees to offset carbon emissions or by investing in new technologies) or by setting up sanction mechanisms for the users.

If we think of the atmosphere as something in which we all have to invest in order to preserve it in an acceptable state, then indeed we can speak of a public good that requires, in turn, collective action. Such a situation in which an individual has incentives to free-ride on the provision of a public good needed to solve a commons dilemma is often referred to in the literature as a *second-order dilemma* or a *higher-order collective action decision* (see e.g. Olivier 1980; Okada 2008).

Referring to climate change, Milinski et al. (2008, p. 2291) introduce the concept of *collective-risk dilemma*, where the main question is: “Will a group of people reach a collective target through individual contributions when everybody would suffer if the group fails to achieve the target?” The target, in this case, is a level of compliance to a certain rule that is high enough to avoid dangerous climate change. In real-life the rule could be anything ranging from monetary contributions to a climate change abatement fund to collectively-agreed limits on CO₂ emissions. Regardless of the kind of target or rule, the main characteristics of a collective-risk dilemma that distinguish it from other situations are:

- “i) people have to make decisions repeatedly before the outcome is evident;
- ii) investments are lost (i.e. no refunds);
- iii) the effective value of the public good (in this case, the prevention of dangerous climate change) is unknown;
- iv) the remaining private good is at stake with a certain probability if the target sum is not collected.” (Milinski et al. 2008, p. 2291)

The idea of a target is particularly important in this discussion, especially since it includes, once again, a subjective component: even if humanity agreed that avoiding dangerous climate change is the collectively desired outcome, there would still be the question of what constitutes dangerous (i.e. an acceptable level of risk), as well as what kind of specific targets should be set and how they should be achieved. Because in reality preference coordination is highly unlikely, it makes sense for our experiment later on to design a decision situation that would dynamically allow people themselves to choose both the strategy and their target.

4.3. Behaviour in Environmental Commons Dilemmas

4.3.1. The Prisoner’s Dilemma as a Model: Uses and Limitations

When it comes to discussing behaviour in an environmental commons dilemma, it is the prisoner’s dilemma (PD) game that has been often used to model the interaction between different users of the resource. If people were to behave “rationally” as

defined by traditional game theory, i.e. by maximizing their own individual profit, they should always choose to appropriate from the shared resource as much as possible, without any kind of concern regarding its availability to others. A different approach would be to cooperate with other users, so that it is possible to achieve a collective “optimum” or – in game theoretic terms – the Pareto-superior solution.

According to Hardin (1968), who was the first to speak about a *tragedy of the commons*, when there is open access to a CPR individuals will follow the rational choice postulate. Because the resource is finite, this situation – he thought – would necessarily lead to overusing and eventually the destruction of the resource. Hence, the solution is to either privatize or regulate resource use by government control. A similar kind of situation can also exist with public goods – if contribution to the provision of public good is not enforced, presumably nobody would invest in maintaining that good which will in turn never emerge, will deteriorate or even cease to exist.

Critiques of the external authority theory that emerged from Hardin’s article have argued, though, that other solutions exist as well. For instance, Ostrom et al. (1999) found empirical evidence that the creation of institutions and self-imposed norms have long functioned as mechanisms for sustainable management of CPRs and PGs.

Nevertheless, the apparent paradox of rationality involved in dealing with environmental commons has inspired numerous CPR and PG experiments. A typical public good game requires players to invest a share of their endowment in a collective asset. While the payoff maximizing strategy is to contribute nothing, collectively the players would be better off if they all invested. Depending on the conditions inherent to the game, it turns out that indeed people will deviate from the rational choice norm. Some of the more general conclusions of such experiments are that²⁰:

- subjects cooperate in one-shot PD games about half of the time and contribute about half their endowments in public goods game (Sally 1995; Ledyard 1995 qtd. in Camerer 2003);
- low rates of voluntary cooperation can be remedied by institutional arrangements or informal mechanisms (Ostrom 2000 qtd. in Camerer 2003);
- when the games are repeated with random “stranger” rematching, cooperation and contribution rates decline over time (Fehr and Gächter 2000c qtd. in Camerer 2003); punishment, however, is very effective even when it is costly, raising contributions to more than half of the endowments (Yamagishi 1986 and Fehr and Gächter 2000c qtd. in Camerer 2003);
- in CPR situations, communication and the ability of users to define own agreements and punishments leads to near optimal rates of resource use (Cárdenas and Ostrom 2004);
- large group size and non-monetary interests will decrease cooperation (Sally 1995);

²⁰ The first three observations are based on an overview given by Camerer 2003, pp. 45-48.

- an increase in the relative level of temptation²¹ will lead to a decrease of the observed level of cooperation (Ahn et al. 2001);
- repeated interactions with the same individuals will lead to more sustainable resource use (Ahn et al. 2001).

Of course, the findings outlined here provide only a limited image of the kind of behaviour that people will exhibit in environmental commons dilemmas. Many different experiments have been carried out, each presenting various types of deviations from the traditional idea of rationality. The common point of all these studies is, though, that when it comes to protecting a CPR or to providing a PG, cooperation *does* emerge. Consequently, the challenge is to understand the triggers and the conditions under which it happens. This is what we will try to briefly look at in the following section.

4.3.2. Cooperation: Explanatory Models

With regard to the factors that lead to cooperation in PD games, many efforts have been made in an attempt to explain the observed human behaviour. Here we mention a few explanatory models that are encountered most often in the literature. Most of these follow the lines of what Camerer (2004) called “a theory of social utility” (see Section 3.3.2).

Equity, Equality and Fairness

One of the most straightforward explanations of cooperation refers to a *payoff inequity aversion*. Fehr and Schmidt (1999) argue that the deviations from the rational choice postulate that are observed in some situations can be explained by the presence of a proportion of people that are “fair-minded”, i.e. they care more about equal and/or fair outcomes than about their own selfish interest. These people incorporate in their utility functions the relative payoffs of all the participants in the interaction, perceiving a PD game as a coordination one. Depending on the decision environment and the proportion of such people, different outcomes can occur. For instance, experiments have shown that a minority of “fair-minded” people could determine even the selfish ones to cooperate fully in a PD game with punishment (Fehr and Schmidt 1999).

Other research along the same lines has discussed the link between inequity aversion and social value orientations. The theory of social value orientations assumes that individuals are pursuing different goals when confronted with a decision and identifies typically the following categories (see Eek and Gärling 2008): *cooperators* (or *prosocials*) – those who try to maximize the joint welfare, *competitors* – people interested in their relative gain to others, and *individualists* – those maximizing their absolute gain (the last two types also referred to as *proselfs*). In an experiment that, we believe, supports the theory of inequity aversion, Eek and Gärling (2008) prove that prosocials are actually more interested in obtaining equal outcomes than in

²¹ According to the definition given by Ahn et al. 2001 and within a PD game matrix, *temptation* T is understood to be the payoff for playing “defect” when the other plays “cooperate”. Other similar terms include: *reward*, R , the payoff when both players choose to “cooperate”, *punishment*, P , when they both choose to “defect” and *sucker*, S , the payoff of playing “cooperate” when the other plays “defect”, with $T > R > P > S$.

maximizing the joint outcome: they preferred more an outcome of 500 for self and 500 for others than 500 for self and 800 for other.

Reputation

Other explanations of cooperation bring into discussion reputation. Knez and Camerer (2000), for example, study how reputations built by players in some games and a history of efficiency will transfer to other games, even though the new payoff structure no longer induces cooperation, as in the case of a PD game. Similarly, Ahn et al. 2001 prove that a history of play in coordination games where cooperation has been achieved will lead to increased cooperation in future PD games, and also that the rate of cooperation will be significantly higher in the fixed-matching protocol than in the random one (i.e. when people play against the same opponents every time). In other words, people will cooperate more when they interact repeatedly with the same individuals and reputations can be built.

Reputation-building is not, however, only engaged in with view to higher benefits later in the future. Experiments conducted by Andreoni and Miller (1993) suggest that some people build reputations as cooperators not because they just want to appear altruistic and defect later, but because they truly are so. For a stable fraction of the subjects, additional utility seems to be derived from cooperation.

Information & Context of the Game

A third category of explanations to why people cooperate in commons dilemmas refers to the information that they have available and the context of the decision situation.

According to Cárdenas and Ostrom (2004) people use a three-layer information model in order to make decisions. More specifically, they look at the material payoffs – their actual gains, the group-context – whether others can be trusted or not, and their identity – whether a certain behaviour would be consistent with their set of values and image of self. As Cárdenas and Ostrom (2004), players will try to complement the information about the basic structure of the game with information about the other layers. When the situation is of such sort that it is difficult to gather information about the other players and the group-context layer is thus useless, cooperation is less likely to emerge. As Bowles (qtd. in Cárdenas and Ostrom 2004, p.315) suggests: “the more the experimental situation approximates a competitive [...] market with anonymous buyers and sellers, the less other regarding behaviour will be observed”.

The role of the different information layers can also be observed in CPR experiments where face to face interaction between participants is allowed. When people are able to communicate, trust-based agreements can be settled which in turn will lead to successful cooperation (Ostrom 2006).

Critique

So far we have presented a few theories that could explain cooperation. Although this list is not at all complete, the value of these models lie in the fact that they give an overview of the multitude of variables that might influence the decision to cooperate, as well as of the relationships between them. The problem with these models is, though, that they each focus on certain characteristics of the users or of the context, while disregarding others.

If, let's say, we regarded fairness considerations as the most important for cooperation, the next challenge would be to find out what is considered fair and what not. Research in this direction has shown that people think it is fair that everyone contributes to the provision of a public good in proportion to one's assets (see Kazemi and Eek 2008). However, in situations of asymmetric resource dilemmas – i.e. when there is unequal access to the resource – individuals have difficulties in defining what is fair (Kazemi and Eek 2008). Similarly, as we have already mentioned earlier, equality functions as a rule only in some limited circumstances.

The same applies for reputation and information gathering, while the emergence of norms is not sufficient to ensure cooperation, but it also needs to be backed up by mechanisms for negotiation, monitoring and rule enforcement (Ostrom et al. 1999).

Instead on concentrating on limited, one-sided, explanations of cooperation, it makes sense to take a more holistic perspective and talk about *factors* that *enhance* cooperation in commons dilemmas. If our unit of analysis is the entire context of the decision situation, then mapping out all the variables that could lead to cooperation might be worthwhile. Ostrom (2007) suggests a framework that contains a broad range of variables that could be used to characterize the socio-economic setting and to track the influences these might have on cooperation. However, it is not our scope here to elucidate the fine details of such a variable map. In our study the unit of analysis is the individual and the focus is on decision-making. Therefore, in the next section we will discuss a decision-making model that has been particularly developed to explain behaviour in environmental risk situations and that we find to be more robust than others.

4.3.3. Protection Motivation Theory

From an individual perspective, some of the considerations mentioned above could well fit into a multiple-motive model of decision-making. Originating from psychology, the protection motivation theory attempts to elucidate the key elements that are involved in valuating the trade-offs between risks and benefits that occur in relation to environmental resources.

According to this model, as presented by Steg and Vlek (2009), every risk situation is approached by performing a *threat appraisal* and a *coping appraisal*. The threat appraisal component deals with estimating the risk situation based on one's needs and values, while the coping appraisal component refers to the general efficacy of an imagined action as perceived against personal knowledge, skills and support. These two sides of the evaluation in turn determine the temptation to accept the risk and continue to reap the benefits, or respectively the motivation to reduce the risk at a certain cost. This balance between benefits/costs, or temptation and motivation to act will materialize into an intention (protection motivation), and will consequently determine the behavioural strategy to be adopted. The schematic representation of the protection motivation theory is given in Figure 12.

The beauty of such a model lies in the fact that it allows not only for a rationalization of the actions to be carried out, but also for incorporating needs and values, and any other kind of perceived benefits, be they intrinsic or extrinsic. Also, this model is consistent with the finding that issues such as reputation or fairness are incorporated in the utility functions of people. While it is not easy to test this model as rigorously as other explanations, more suitable for economic experiments, we think it is however a useful behavioural framework of how decisions are made, especially

since it brings into discussion a new component besides risk, i.e. the coping appraisal. Coping appraisal refers to evaluating the situation in terms of one’s own abilities and results into reframing the situation in such a way that it can justify action or inaction (see Figure 12).

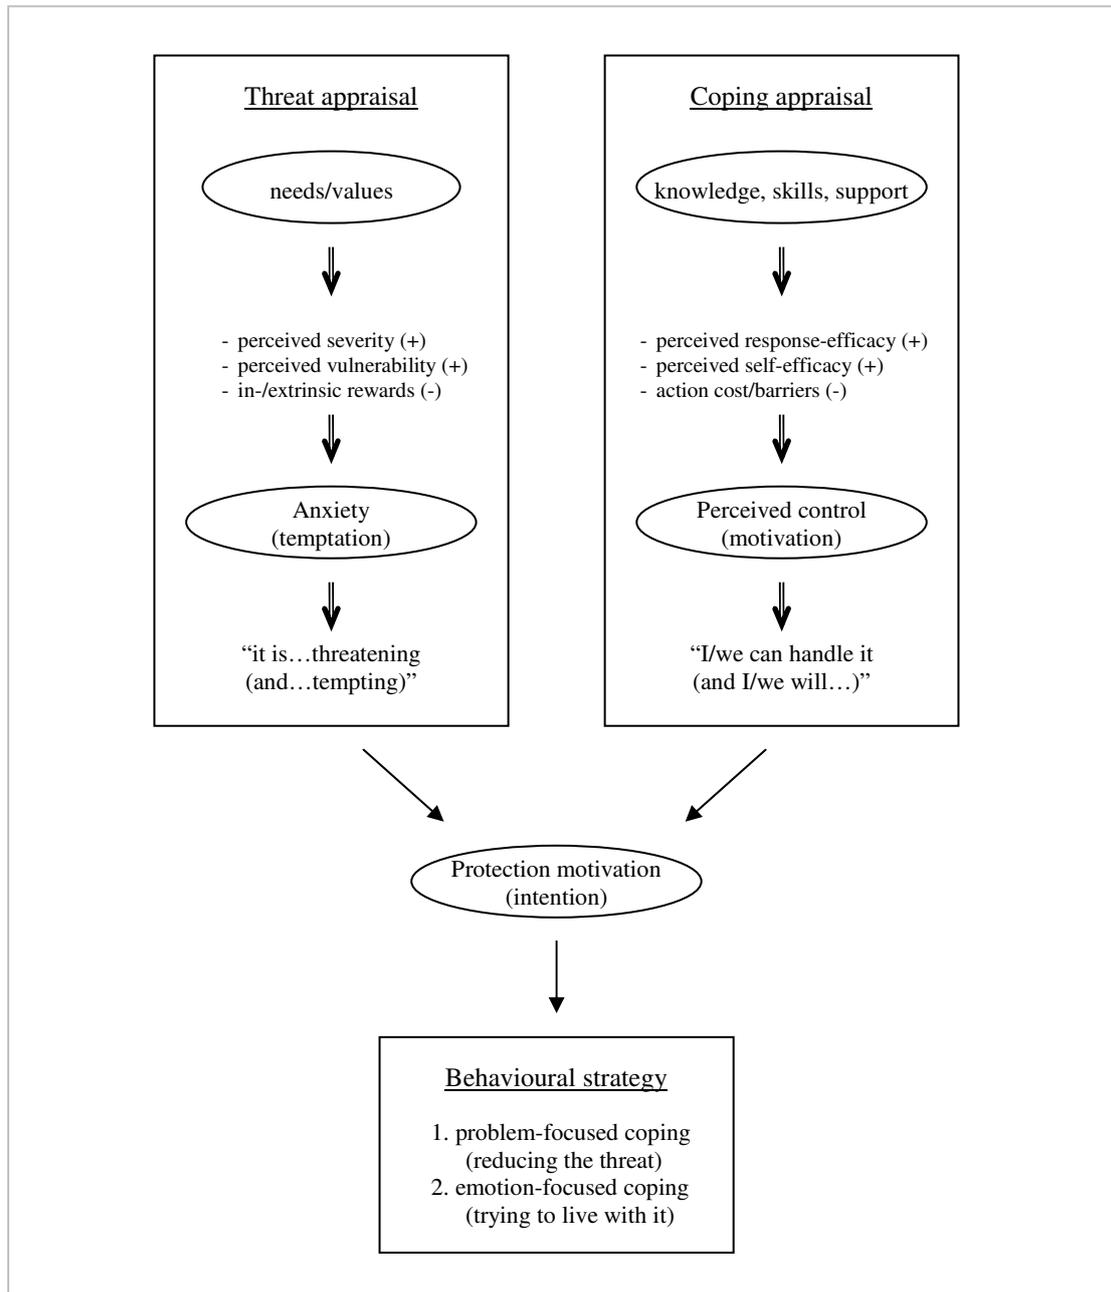


Figure 12. Schematic Representation of Protection Motivation Theory (reproduced from Steg and Vlek 2009, p. 108).

Perceived Risks

An important issue to mention here is that, in the case of such models, risk is no longer defined simply as an “objective” probability, but it is rather a matter of intuitive judgment (Raaijmakers et al. 2008). Indeed, Steg and Vlek (2009) also acknowledge that there are many dimensions of perceived risks, among which they collected 11 from different empirical studies (see Box 3). Nevertheless, an interesting risk configuration model is provided by Raaijmakers et al. (2008) and includes three

components that, in our view, are also embedded in other studies as it will be shown later.

1. Potential degree of harm or fatality
2. Physical extent of damage (area affected)
3. Social extent of damage (number of people involved)
4. Time distribution of damage (immediate and/or delayed effects)
5. Probability of undesired consequences
6. Controllability (by self or trusted expert) of consequences
7. Experience with, familiarity, conceivability of consequences
8. Voluntariness of exposure (freedom of choice)
9. Clarity and importance of expected benefits
10. Social distribution of risks and benefits
11. Intentionality of harm

Box 3. Dimensions of perceived riskiness (based on Steg and Vlek 2009, p. 109)

While attempting to study flood risk perceptions, the authors propose three characteristics as being relevant to the understanding of risks (of hazards) and human responses to them: *awareness*, *worry* and *preparedness*. First, *awareness* refers to the knowledge of the risk among those who are exposed to it. A distinction is made between three levels of awareness: *expert awareness* (least uncertainties about probabilities and consequences), *underestimation* of probabilities of occurrence or consequences, and *ignorance* in regards to risk exposure. Second, *worry* describes the extent to which people dread the hazard and depends on their expected severity of the consequences. Third, *preparedness* defines the ability to cope with a hazard once it happens and to recover from it ex-post. According to Raaijmakers et al. (2008) more awareness increases worry which, in turn, will lead to better preparedness. On the other hand, high preparedness will reduce worry – as people start to feel safe – and, over longer periods of time, this might also lead to lower awareness, to the extent to which the hazard risk is less discussed (Figure 13).

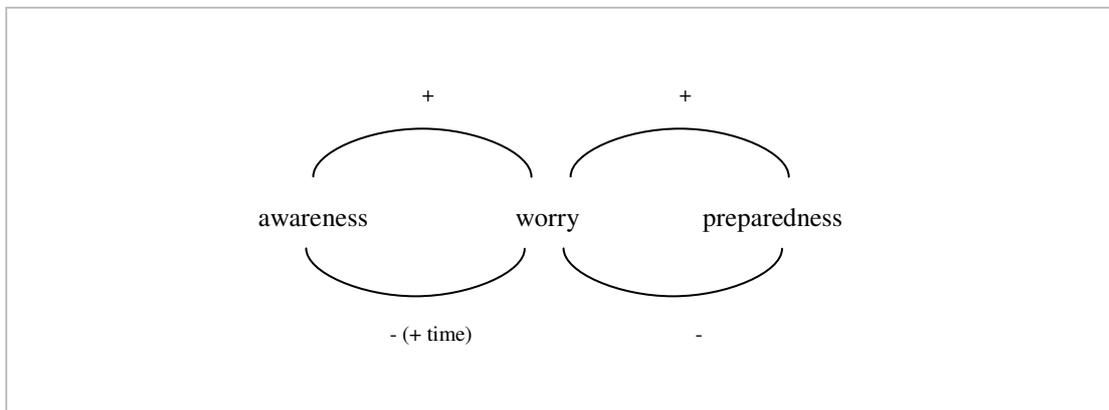


Figure 13. Relationship between (Flood) Risk Characteristics (reproduced from Raaijmakers et al. 2008, p. 312)

While the Raaijmakers et al. (2008) model is particularized for dealing with flood risks, which is a different situation from the one posed by climate change, it is mentioned here because of its relative simplicity. The concepts involved as well as some of the relationships described are often found in other articles, albeit under different names.

4.4. Lessons from Previous Climate Change Experiments

We have now sketched some of the major findings related to behaviour in environmental commons dilemmas in general, and we presented a useful framework for decision-making that accommodates for these findings. To what extent are these lessons applicable to climate change?

Few studies have attempted to examine people's behaviour in social dilemmas framed as climate change. This is however important because in real life it is difficult to disentangle the cooperation aspect of this problem from the risk, time frame or uncertainty issues.

First, when it comes to observing behaviour in climate-change-like situations, the study of Milinski et al. (2008) is the only one that we are aware of at this moment. In an economic experiment they study exactly this relationship between (collective) risk and coordination. More precisely, they set up a public goods game where several subjects have to invest a certain amount of their endowment into a common fund, with the objective of collectively reaching a certain target sum of 120€ in the fund. Failure to reach the target will determine the computer to "throw dice" and thus randomly decide, with a probability p , if the subjects will lose all their savings (personal savings = total personal endowments - total personal investment) or not. They test several values for p – 90%, 50% and 10% chance of losing the savings – and observe that in the high risk treatment 5 out of 10 groups successfully manage to collect the target sum. Moreover, some "irrational" behaviour is seen in some of the other treatments when a few subjects continue to sub-invest in the last rounds although it becomes evident that there is no way to achieve the target sum unless a large number of players invest. As hypothesized by the authors, this might be due to a reluctance to reward behaviour that is perceived as unfair.

As far as the factors that influence behaviour in climate change issues are concerned, Milinski et al. (2006) investigate the effects of reputation and information on the willingness to pay for sustaining the climate. In a public goods game, instead of dividing the common resource among the participants it is promised that the pool will be used for real-life climate change related actions (invested to encourage people to reduce their fossil fuel use). Not surprisingly, users do behave altruistically in the control treatment, and will contribute even more of their endowment in non-anonymous rounds and when they are better informed about the consequences of climate change. This indicates that policy makers should strive to make scientific certainties explicit to the large public, as well as to design "strategies to improve the social reputation of people investing in climate change" (Milinski et al. 2006).

With regard to the role of the information in shaping behaviour, Sterman and Sweeney (2007) point out to the need of ensuring that the public understands well the system dynamics behind issues such as climate change. In a study, they test the understanding of basic stocks and flows dynamics on a group of MIT graduate students. The subjects are provided with graphic information describing the relationships between historical greenhouse gas (GHG) emissions, past and future atmospheric concentrations. They are then asked to sketch the future emissions trajectory required to stabilize CO₂ concentrations as indicated by the given concentration graph. This also requires, from their part, indicating the estimated net CO₂ removal rate. Most respondents fail to understand that a stabilized future concentration requires immediate drop of emissions below the net removal rate. Instead, 84% exhibit a pattern-matching behaviour, drawing an emission line that

follows the shape of the concentration one, but is above the removal rate, as if they believed that simply stopping the growth of emissions will stop the growth of the GHG concentrations. This belief violates the mass conservation principles and suggests that the general public does not understand the dynamics between emissions, concentrations and temperature, nor the high inertia of the climate system.

From a cognitive perspective, information is important as it affects risk perception (and thus subjective probability assessments), which in turn constitutes an essential component for solving the benefit-risk dilemma. As indicated by a multi-agent model simulating behaviour in climate change situations, learning is a key process here, especially in a world of imperfect information where the mix of policies and aspirations is forever changing (Janssen and de Vries 1998). When it comes to assessing the expected value of temperature increase, for instance, social studies confirm that knowledge is generally updated in an approximately Bayesian manner. (Cameron 2005). However, in order to have an actual effect on behaviour, information needs to be reliable: Cameron (2005) emphasizes that disagreements among external sources of information will decrease the amount of attention paid to them.

We have now moved into the field of (psycho-)sociological research. Because of the very nature of the problem, risk perception seems to be a favourite variable when it comes to climate change related research, but only in a few cases is it studied as an independent variable affecting behaviour. O'Connor et al. (1999) hypothesize that willingness to support personal and governmental behaviours is a partial function of: general environmental beliefs and risk perceptions. A major finding is that perceptions specific to climate change and general environmental beliefs are independent as accounts for behavioural intentions. When it comes to understanding even more about the dimensions of risk perception, knowledge about the causes of climate change is, once again, found to be an important predictor of behavioural intentions.

Continuing with the (psycho-)sociological research, a study by Viscusi and Zeckhauser (2006) seems to be particularly relevant, as it also brings into discussion the decision theory premises mentioned in Chapter 2. In a complex assessment of the perception and valuation of risks of climate change, Viscusi and Zeckhauser (2006) found that many of the participants' answers elicit a blend of rational and behavioural decision making:

“In the climate-change arena, behavioral decision tendencies are like a fun-house mirror: They magnify some estimates and shrink others, but the contours of rational decision remain recognizable” (Viscusi and Zeckhauser 2006, p. 168).

Among others, they show that information about climate change correlates positively with increased temperature estimates, which confirms the awareness-worry positive link in the model of Raaijmakers et al. 2008. Also, when subjects make higher temperature increase estimates they are also more in favour for aggressive policies, which could justify why worry triggers more preparedness. However, if we understand preparedness as the effect of policies, we should observe that, at least in the case of climate change, inferring a higher ability to cope with potential hazards simply from higher worry means to rush towards a wrong conclusion. In other words, the fact that people worry more does not necessarily mean that they will be more able or willing to actually cope with climate change related challenges. As indicated by

Leiserowitz (2006), there is a contradiction between risk perceptions (understood here as worry) and policy preferences, that engage people in a “wishful thinking” state of mind: while they support climate policies at national and international level, they oppose policies that directly affect them, such as taxation. This suggests that worry might have some intentional effects, but the latter do not necessarily translate into observable behaviour.

4.5. Conclusions

In this chapter we tried to provide an overview of some of the concepts that could be useful in understanding the problem of climate change and to explain how these are linked to individual economic behaviour.

We discussed environmental commons dilemmas and empirically-proven ways to overcome them, but we also showed that climate change is a special type of decision situation due to its global, yet at the same time local, character. First of all, the atmosphere shares the characteristics of a global commons, where issues such as cultural diversity, scale and requirements for unanimous agreements, to name only a few, constitute great challenges for any attempt at collectively cooperation (Ostrom et al. 1999). Secondly, those who contribute to climate change are not necessarily those who will be the first to bear the effects. Thirdly, the fact that actions taken in the present only have consequences later in the future constitutes an incentive to preserve the status quo, especially because people have been shown to have a hyperbolic time preference (i.e. they value short-term benefits more than long-term ones). Lastly, the uncertainties associated with climate change might make people even more reluctant to invest time and effort into addressing what is already a complicated matter.

When it comes to understanding how people might judge and deal with complex issues, we concluded from previous studies that it is in the realm of psycho-sociological research that we might find some answers. Consequently, we presented some theories and findings that deal with motivations, perceptions and attitudes. Perhaps one of the most important lessons of this chapter remains that we cannot isolate people from their ideas and interpretations about the world. Hence, we will try to integrate also some of these dimensions in our conceptual model of decision-making under risk and uncertainty.

5. Modelling the Climate Change Decision Context

5.1. A View on the World

After giving a broad overview of the literature on decision-making and climate change as a decision situation, the time has come to integrate these findings into an interdisciplinary behavioural model that is particularly tailored to accommodate the unique characteristics of climate change as a policy problem.

However, designing a decision situation which resembles that of climate change policy involves, first of all, making several assumptions about the world. Therefore, before we move on to the conceptual model itself, we briefly present here how we see the relationship between individuals and the surrounding world, especially when it comes to making decisions.

We take a philosophical stance by assuming that there is an objective reality that is only accessible to us by means of individual perceptions. This objective reality is the sum of nature, societal and individual characteristics. For every decision problem, the interactions between the three spheres are not known to us directly, but only as individual and subjective evaluations, i.e. perceptions. These perceptions are shaped by different degrees of knowledge about the objective reality and they ultimately lead to actions. In Figure 14 we represent the interplay of objectivity and subjectivity in decision problems. While not shown in the figure, actions may have an influence on the objective reality through any of the three spheres.

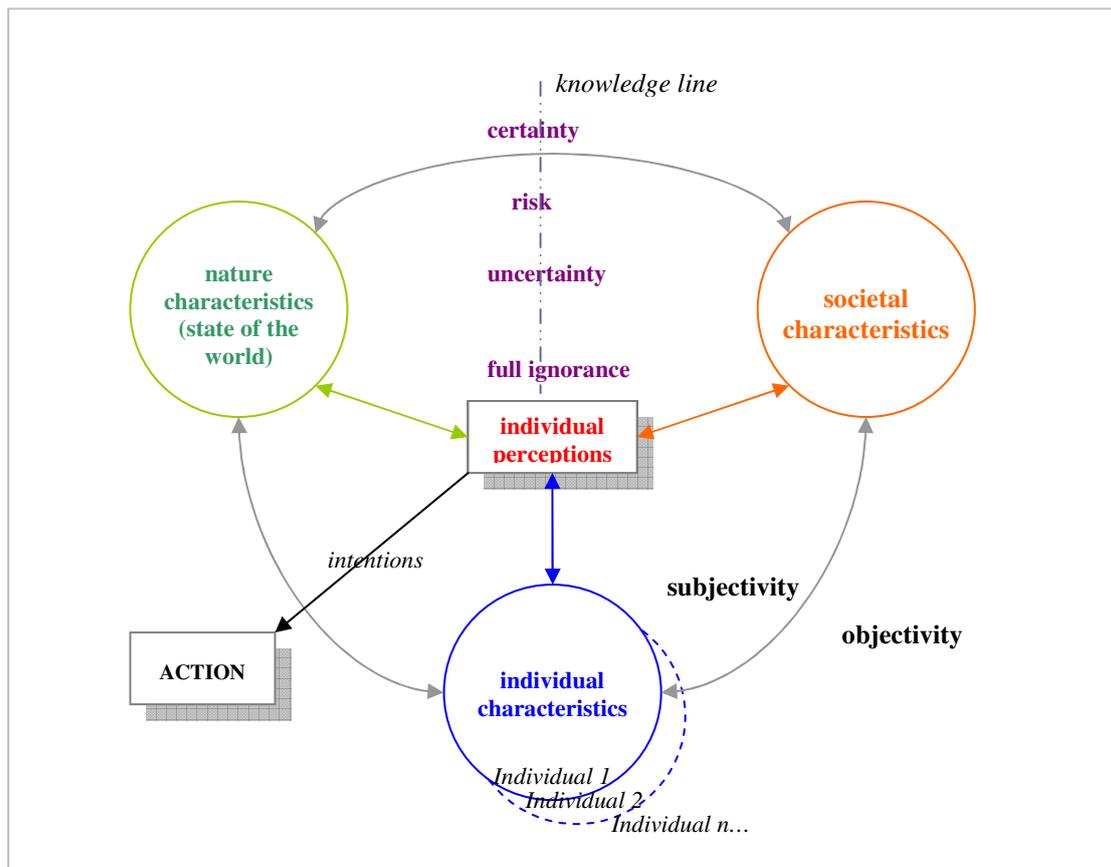


Figure 14. Decision Problem: Objectivity and Subjectivity

The “knowledge line” refers to the information that we have about how the world functions and we place here some of the concepts that we presented in Chapter 2 and that were shown to be relevant for decision-making. The closer to objectivity our information is, the higher the degree of knowledge. Since, in an absolute sense, we believe that nothing can ever be known for sure, “certainty” is placed on the border between subjectivity and objectivity.

A decision problem can concern one or more individuals and despite the possible differences in perceptions, common actions are possible through information exchange mechanisms. These processes are not discussed here as our main focus is on decision-makers at an individual level. Nevertheless, it is important to acknowledge that this view on the world does allow for collective action, thus also being applicable to decision problems – such as climate change – that concern multiple actors.

5.2. Conceptual Model

Having specified the paradigm in which we operate, we can now present a conceptual model particularized to the climate change decision situation that will be useful for studying economic behaviour under risk and uncertainty. We construct this model through a grounded-theory approach, where we strive to integrate as much as possible the three spheres described above, by including variables from all domains.

It is relevant at this point to mention that by no means do we claim that this model is exhaustive in the factors it presents. There are for sure many more variables that intervene in this decision process, which make the research object of sciences ranging from climatology to psychology. The value of our effort is, however, in the interdisciplinary approach that we aim for.

In the simplest form, the mechanism can be described as following (also see Figure 15):

Climate change is expected to affect humanity by certain damage (in the case of this study this will be quantified in monetary terms). We assume that there is an objective *actual damage* that one might incur due to climate change. This damage depends on natural conditions (*problem malignancy*), but also on whether people decide to address the problem of climate change by taking action or ignoring it. From the range of possible options, we focus here on *individual economic behaviour*: one can decide to reduce the threat of climate change at a cost, to ignore the problem and focus on his short-term benefits, or to adopt a strategy that is somewhere in between these two extremes. According to the protection motivation theory (see Steg and Vlek 2009), a behavioural strategy is chosen following an internal evaluation of the risks involved and the perceived control. In our model, we call these two elements *external risk perception* and, respectively, *solvability perception*. At any point t in time, the *external-risk perception* is a subjective interpretation of the current *actual damage*, and it is shaped, among other things, by *uncertainty* and the *individual and collective risk discourses* that dominate the decision situation. The *collective risk discourse* refers to the main societal message regarding how individual contribution can reduce the collective risk of climate change, while the *individual risk discourse* encompasses the mainstream view about whether climate change will affect all people equally or not (see the “Variables” section for further clarification). The *solvability perception* variable depends, in its turn, on the *degree of control over risk*. The latter is interpreted here as being a (fixed) characteristic of the decision situation. On the other hand, *solvability perception*, like *risk perception*, is strictly subjective

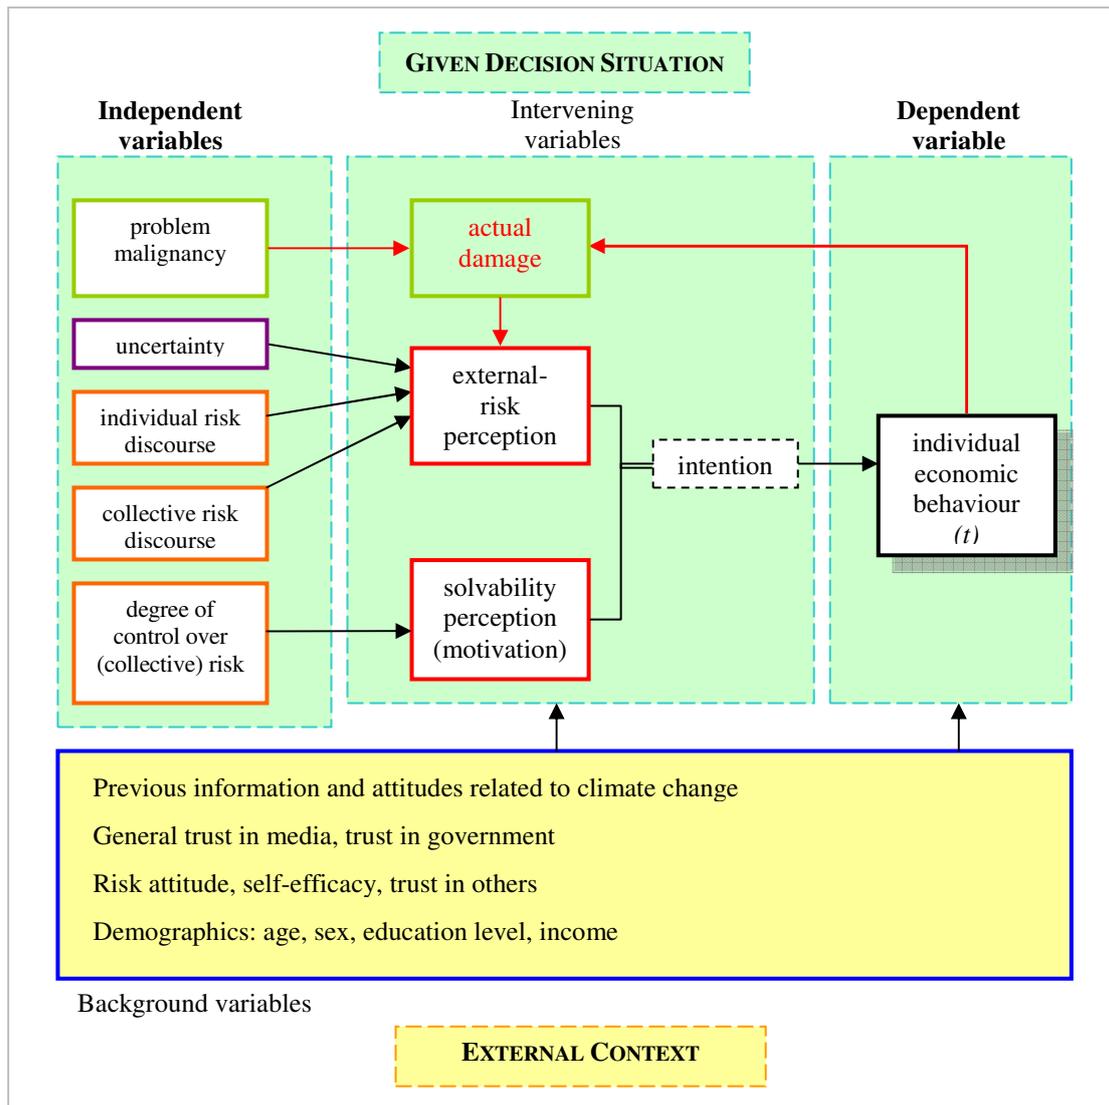


Figure 15. Conceptual Model

The mechanism above describes the factors that lead to a certain economic behaviour within the limited decision context of climate change. Note that the border of each variable respects the colour convention from Figure 14, in order to illustrate its belonging to one of the three reality spheres (nature, society, individual).

From the point of view of the decision-maker, the “independent variables” are parameters that characterize the decision situation and are relatively fixed, while the “intervening variables” box consists of those factors that are constantly updated based on new information/feedback from the environment. However, when pondering their actions, people cannot completely isolate themselves from their identity and their broader system of beliefs. When making evaluations about the decision situation they are in, they also involve their values, or make use of the experiences and lessons from other aspects of their lives. It is this set of factors that we group in the *background variables* box.

Consistently with the finding of Viscusi and Zeckhauser (2006) that when it comes to climate change people’s appraisal is a combination of rational and behavioural elements, our conceptual model tries to integrate both elements from decision theory and behavioural science. The “given decision situation” part of the

model lists the variables that are directly linked to the climate change problem and decision-making (in green), while the “external context” part comprises the collection of prior/external characteristics or beliefs that the individuals bring as additional input to their decisions (in yellow).

5.2.1. Variables

We now explain in more detail the variables in our model. In the next subchapter, we will discuss the laboratory experiment as our chosen method to study this model and a formalization of these variables will follow.

Dependent variable

Individual economic behaviour – quantitatively, we are interested to see how “rational” (in Bayesian terms) economic behaviour of individuals is, under different conditions. The underlying assumption is that there is an “optimal” behaviour, implicit in the sustainable development discourse, where we accept short-term costs that are *necessary* to sustain the current level of needs satisfaction in the future, but nothing more (i.e. we still develop). Qualitatively, we are also interested to observe the behavioural strategies that people develop when coping with different characteristics of the decision situation. If it turns out that people are not “rational”, then we will quantitatively look at the level of costs that they are willing to incur in order to reduce risk.

Independent variables

Problem malignancy – so-called stochastic uncertainty with regard to climate change. We do not exactly know in which ways will climate change and/or ecosystems adapt, although different scenarios of climate evolution can be made. Climate change might “hit” us all of a sudden or, on the contrary, have effects that increase gradually. We can *expect* different levels of malignancy and work with them as with scenarios.

Uncertainty – this refers to the degree in which decision makers have access to the best scientific knowledge about climate change. It is explicated along three dimensions/sub-variables (these were not shown in the picture due to space limitations):

- 1) *Knowledge of the causes/mechanisms of damage*: whether the decision maker has “expert awareness” in the terms of Raaijmakers et al. (2008) about which steps lead to which outcomes. This knowledge also includes awareness on how economic behaviour will further influence the evolution of climate change. This is factual information and is different from the more “subjective” interpretations and beliefs that the decision maker might already have about climate change.
- 2) *Information availability*: whether the decision maker obtains information about the feedback of his actions on the climate system in time to influence his next decision or not. This is consistent with the real situation where consequences of a decision may only become available long after it was made.
- 3) *Information reliability*: whether the feedback that the decision maker receives about the consequences of his actions is accurate or distorted by

other factors in the social environment. These distortions can occur, for instance, as miscommunications between science and public.

Along the lines proposed by Milinski et al. (2008), one should distinguish *collective-risk* from *individual risk*. In the case of climate change, a collective risk could be, for example, the probability that the global mean temperature will increase by 4°C. On the other hand, whether climate change will affect or not one individual's property is a matter of chance and thus, individual risk. Making this distinction between the two types of risks is relevant particularly because people might exhibit different behaviours when risks are close to their daily lives than when they are more indirect, or far away from 'now' and 'here' (temporal and spatial dilemmas).

Collective risk discourse – whether the economic behaviour of the decision maker is understood as contributing to:

- 1) Adaptation, or
- 2) Mitigation/avoidance of climate change.

Individual risk discourse – whether climate change is seen as:

- 1) “One for all and all for one”: a global problem, where the *actual damage* affects us all with equal probability
- 2) “To each his own”: although there is a collective actual damage, there is a chance (probability) that some individuals will not be affected by this damage, depending, for instance, on where in the world they live.

Degree of control over collective risk – whether or not individual actions have a direct effect on the actual damage or not. In real life, the actual damage of climate change depends on the collective actions taken by many individuals. However, in an experimental setting it would be possible to manipulate this variable as to eliminate the social dilemma factor and observe the individual economic behaviour under full control over the collective risk.

Intervening & Background Variables

The intervening variables represent some (cognitive) factors that presumably mediate between stimuli and behaviour. Background variables constitute a collection of prior characteristics and beliefs of the subjects that might influence the entire spectrum of variables within the “given decision situation”. Neither the intervening, nor the background variables can be manipulated.

External-risk perception (worry) is defined here closely to the concept of “worry” defined by Raaijmakers et al. (2008) and is the result of the “threat appraisal” in the language of Steg and Vlek 2009. This is not to be confounded with the collective or the individual risks that are pre-set within the game, but rather to be understood as an attitudinal component that forms during the deliberation process.

Solvability perception (motivation) refers to whether or not the subject feels in control of the collective risk while playing the game and can be interpreted as the result of the “coping appraisal” from the protection motivation theory (Steg and Vlek 2009).

Intention is a variable mediating between perceptions and behaviour. We introduce it mainly to accommodate for Weirich's (2004) distinction between “effect”

and “disposition to act” (also see footnote 5, p. 7), but we will not attempt to measure it rigorously, hence the dotted line around it.

The actual damage constitutes a special type of intervening variable that is dynamically updated at different time intervals. Depending on the individual economic behaviour vis-à-vis climate change, the actual damage could increase or decrease in time. This variable thus captures the evolution of climate change. To the extent to which this evolution has a feedback effect on the social system, it also influences the external risk perception from one moment in time to another.

As for the background variables, most serve as controls for the independent variables and are generally self-explicative, thus I will no longer insist on them here. While no studies have been made about the role of demographic variables in relation to the climate change issue, this information might also be interesting to collect when the study sample is diverse.

The intervening and background variables play an active role in shaping the behavioural strategy; hence they will be important for the qualitative analysis of the dependent variable.

5.2.2. Hypotheses

Based on the conceptual model presented above many relationships between the variables can be tested, but our initial focus is on studying the effects of risk and uncertainty on individual economic behaviour, thus we give priority to the independent variables and use the other variables as complements for a more robust analysis.

Also, as we will see in the next chapter, the formalization of our model introduces different levels for each variable, resulting in a large number of possible combinations. Since practical limitations have restrained us from having a full factorial design in our experiment, not all of the links could be explored. Nevertheless, we still present here some general hypotheses that can be advanced based on this model, although not all of them will be tested in this study.

- Individual economic behaviour is closer to the “optimum” level when expected problem malignancy is higher (climate change effects are less predictable);
- Individual economic behaviour will be closer to the “optimum” level when the uncertainties are low (i.e. information is readily available – no delay – and reliable – no noise);
- Individual economic behaviour will be closer to the “optimum” level when the decision makers have full control over the collective risk (i.e. their decisions have immediate, unmediated effect on the outcomes).

No expectations are yet made about the effect of the risk discourses on the individual economic behaviour, as no such studies have been made before. The effects of different risk discourses can only be explored as scenarios.

6. Testing the Model

6.1. Method

As already mentioned before, we chose to test the model, and especially the relationships under the “given decision situation” heading, by means of a *laboratory experiment (game)*. In addition, the intervening and background variables are measured through a *questionnaire* and primarily serve as control for the variables manipulated in the game. Also, the questionnaire serves as complementary information in the discussion of the results, by providing an understanding of the individual behaviour strategies and the reasons behind these decisions.

6.1.1. Why The Need of a Game?

The choice to address the research question by means of an experiment is closely linked to the *problem context*, i.e. the climate change issue. Climate change related decisions are irreversible and, since we only have one world, reality does not allow us to experiment with different variable settings. By contrast, in our study we create an artificial decision situation that has the advantage of allowing us to always “reset” the system to the initial parameters and examine the individual behavioural outcomes under different conditions of risk and uncertainty. For instance, in reality we can never know “what would have happened if” the circumstances at a given decision point would have been different. In an experiment, it is possible to create these differences in circumstances *ex-ante* and then observe the outcomes for each of them, resulting into possible scenarios for the future. Also, from the many pairs of glasses that one could use to look at the policy problem of climate change we deliberately choose one that will offer us a view of the issues at hand which is at least partly supported by quantitative analysis.

A second consideration regards the *uncertainty* factor that is of particular interest to us. Climate change related uncertainties only become certainties after long periods of time, thus it is difficult to observe how people adapt/react to new information in a real life setting. The controlled environment provided by an experiment has the advantage that it can “unfold” uncertainty (by feeding information to the participants as it becomes available) faster, and at the discretion of the experimenter. In our game, what initially is totally unknown to the participants becomes clearer as they make decisions and get feedback, so that they can thus adjust their strategies accordingly. Formally, this is called a *differential game*, which means that a payoff at a point in time depends on the payoffs and actions taken in previous decision periods (Finus 2001).

While it is true that this approach isolates the individual from the real-world situation with its ambiguities and complexities, with enough attention a decision situation that resembles the original can be modelled quite accurately. Dynamic, adaptive behaviours to uncertainty can then be easily observed in the lab, within a manageable time span.

Last, but not least, our objective is to explore *behaviours* in climate-change-like decision situations, as opposed to attitudes or perceptions which could be well studied through interviews or questionnaires. Behaviour-wise, people could be observed in real-life settings and measurements could be made about their daily choices (e.g. seeing if they choose the most environmentally-friendly alternative from their set of

options). However, when talking about climate change, the diversity and complexity of people and issues involved may often lead to insignificant correlations or a multiplicity of causal factors. A laboratory experiment, on the other hand, permits a more systematic assessment of people's motives and interpretations in relation to their actions. In addition, as demonstrated by many others in the field of behavioural economics (and traditionally in psychology, too), an experimental set-up allows us to study people's behaviour in a controlled situation and thus permits more solid conclusions regarding the influence of different factors.

6.1.2. CLIMEX – An Interactive Decision-Making Environment

For the experimental purposes described here, an interactive decision-making environment – CLIMEX – was developed from scratch as part of this project, with support from the EU Coordinated Action project on Global System Dynamics²². The programming part was realized by Christian Michel, who developed CLIMEX as a PHP-based environment that runs in any browser. The application allows participants to make decisions over a number of rounds, as will be further detailed in the next section.

The value of this interactive platform lies in the fact that it has been designed to be very flexible, allowing the experimenter to manipulate the independent variables in multiple ways, and thus accommodating many variations of the experimental conditions. Moreover, because it is not a software package, it needs no installation and is accessible online, thus permitting the scaling of this kind of economic experiments to very large samples. More details about the functionalities of CLIMEX, as well as screenshots, are given in Appendix A.

6.1.3. Main Game Setup Description

The main principles of the kind of game that can be defined in CLIMEX are as following:

A participant plays T rounds of a single (or multiple) player investment game. The number of rounds is known in advance. Every round $t=1..T$ he receives an income I ($I=\text{constant}$) which he has to allocate to two funds: (i) an amount $s(t)$ into a savings fund in which the participant can store his income (and which is the basis of his later payout) and (ii) an amount $I-s(t)$ into a climate change insurance fund that mitigates possible future *individual* risks to his savings fund (corresponding to the *actual damage* in our conceptual model). Participants are told that not investing into the climate change fund will forfeit a large part of the savings fund and investing more into it will reduce the risk of losses of the savings fund. In the case of our particular experiment, the instructions that the participants receive contain explicit references to climate change as a decision context (see Appendix B). However, this is a choice that we make in the context of this study, as CLIMEX can also be used with decision situations that are not contextualized, but involve other public good investments.

In every round, before deciding on the investment allocations, participants are informed about expected losses to their savings fund (actual or perceived damage, see

²² <http://www.globalsystemdynamics.eu>

Uncertainty variable in Section 5.2.1.). This information is not always completely describing the objective reality, but is sometimes incomplete, to reflect *uncertainty*.

The “objective” reality is described by the *actual damage* variable, and is calculated in our game as a function $f(t)$ of the accumulated points at time t in the climate change insurance fund. The more points one decides to invest in the fund, the less the damage he might incur to his savings.

It is straightforward to show that (i) from a mathematically rational point of view there is an optimal savings level which is completely determined by the function f and (ii) that given the savings level the time distributions of the savings investments does not matter because only the fund content in the last round is used to calculate the actual incurred damage (i.e. saving much in the beginning and putting all investments into the climate fund at the end is equivalent to fully investing into the climate fund in the beginning and saving all income in the end).

6.2. Formalizing the Model

The experimental approach is pragmatic and consistent with the novel idea of manipulating some variables that are usually only elicited by means of social science methods. The attempt is to have a *factorial design* where every independent variable can take different values and a *treatment* represents a combination of these values.

The dependent and independent variables are presented in Table 2 and their formalization is discussed in the following sections.

Table 2. Variable Formalization

Dependent variable: Individual investment behaviour	
Formalization depends on the game	
Independent variables (with factorial levels):	
1. Problem malignancy	
Linear (<i>L</i>)	Nonlinear (hyperbolic) (<i>NL</i>)
2. Uncertainty	
2.1. Knowledge of the causes/mechanisms of damage	
Function known (<i>K</i>)	Function unknown (<i>U</i>)
2.2. Information availability	
No delay (<i>ND</i>)	Delay (<i>D</i>)
2.3. Information reliability	
No noise (<i>NN</i>)	Noise (<i>N</i>)
3. Collective risk discourse	
Adaptation: damage = fraction of loss	Mitigation/Avoidance: damage = probability of loss
4. Individual risk discourse	
“One for all and all for one”: one die for all	“To each his own”: one die for each
5. Degree of control over collective risk	
Full (single-player game)	Shared (multi-player game)

6.2.1. Dependent variable: Individual investment behaviour

The formalization of this variable depends on those of the risk discourse variables – what we call *risk modes* (see below, also see Risk Modes in Appendix C). In the simplest form of the game, individual investment behaviour can be formalized as the *game score obtained* by the subjects under different conditions. Under the assumption of rationality in the form of utility maximization, looking at the game score would allow for comparisons of performance between different treatments.

However, it is risk modes that define the rules of game and the calculation of the score. In some cases the outcome of any round is entirely chance dependent (participants gain either all their savings or nothing). For such situations other formalizations of this variable are more appropriate, such as the *level of investment in the climate fund*.

In addition, for all treatments in our game it is possible to calculate the “optimal” investment strategy corresponding to a maximum score (or maximum expected value) that could be obtained by a player (also see Appendix B for information on scoring). While in some conditions participants lack the necessary information to derive this “optimal” strategy, we want to see how far away from it their guesses are. Hence, we will also calculate the *deviation* of the individual results from this best possible score. The smaller the deviation, the more “rational” (in Bayesian terms) the investment behaviour is.

In the case of our experiment, we give all the details about the formalization and measuring of the dependent variable in the Section 7.1.1.

6.2.2. Independent variables

a) Problem malignancy is formalized as the *shape and parameters of the damage function f* . This function characterizes the relationship between actions and the actual damage and is an expression of the “objective” risks of climate change, directly affecting the payouts of the players. In this study we choose to experiment with two types of *damage functions* (linear/non-linear), which show different responses of the “climate” to human activity over time. Due to our parameterization, the linear function provides participants with a more difficult problem, as it requires greater efforts (investments) in order to reduce the damage level. That’s why, in this case, problem malignancy in this case is considered to be higher.

1. Problem malignancy	
Linear (<i>L</i>)	Nonlinear (hyperbolic) (<i>NL</i>)

b) The three dimensions of uncertainty are formalized as following:

- 1) *Knowledge of the causes/mechanisms of damage* (function known/not known): whether the damage function is known or not to the players. Under the “function known” setting, participants will receive an additional instructions sheet with information about the relationship between the number of points in the climate change fund and the value of the actual damage at moment t (see Appendix D).
- 2) *Information availability* (delay/no delay): whether there is a time delay or not in the feedback given to participants during the game, from one round to another. When there is a delay, the participant will only find out later what the impact of his decision from a previous round was.
- 3) *Information reliability* (noise/no noise): this refers to a randomness factor that is included in the feedback information given to the participants. The actual damage at point t in time is distorted randomly with a certain noise factor.

In a full factorial design, uncertainty would have 8 factorial levels, as resulting from all possible combinations of the values of the three sub-variables above.

2. Uncertainty	
2.1. Knowledge of the causes/mechanisms of damage	
Function known (<i>K</i>)	Function unknown (<i>U</i>)
2.2. Information availability	
No delay (<i>ND</i>)	Delay (<i>D</i>)
2.3. Information reliability	
No noise (<i>NN</i>)	Noise (<i>N</i>)

c) Collective and individual risk discourses

These variables are manipulated together and operationalized as different sets of rules of the game, as follows:

Collective risk discourses

- 1) Adaptation: the output of the damage function represents the *fraction* of the personal assets to be lost;
- 2) Mitigation/avoidance of climate change: the damage is a *probability* of loss of (all) personal assets at the end of the game.

These risk discourses refer to game settings, and the relationship to real world is our own interpretation. The adaptation mode corresponds to a situation where people are incurring costs as a fraction of their income. In other words, the more you put in the fund the fewer saving you will lose. Investing in the fund is similar to an insurance; investing nothing leads to a significant loss of savings (in percentage), thus one needs to invest something to minimize the losses as possible. On the other hand, mitigation refers to minimizing *the risk* that climate change will happen at all. In the game, this translates into a condition where investing in the fund will decrease the chances that all the savings will be lost. The matter is not of how much one will lose, but of what the chances of loss are.

Individual risk discourses:

- 3) “One for all and all for one”: this option implies that a die is thrown once for all and individual loss at the end of the game is the same for everyone and depends directly on the damage;
- 4) “To each his own”: it is also possible to throw a die independently for every player in which case individual loss at the end of the game is different across subjects.

3. Collective risk discourse	
Adaptation: damage = fraction of loss	Mitigation/Avoidance: damage = probability of loss
4. Individual risk discourse	
“One for all and all for one”: one die for all	“To each his own”: one die for each

Within the game, we call the combination of the collective and individual risk discourses *risk modes*. The rules of the game are different for every risk mode. The exact technicalities can be found in Appendix C.

d) Last, but not least, the *degree of control over collective risk* (full/shared) is formalized as single vs. multi-player game. Under full control, the actual damage and the information reported to the participant depends solely on his investment in the climate change fund. Under shared control the game becomes a social dilemma, as the actual damage depends on the collective contribution to the fund of all players.

5. Degree of control over collective risk	
Full (single-player game)	Shared (multi-player game)

6.3. Experimental Design

6.3.1. Treatments

Putting together all the independent variables explicated before (including the three dimensions of uncertainty), a full factorial design would include 2^7 treatments. Given the time and budget limitations of this study this would have been impossible to complete. Thus, some choices had to be made. The question occurred, of course, which combinations of factors to give up first?

Due to the fact that the multi-player version of the interactive platform was not yet ready at the time we were conducting the experiments, we first reduced the levels of the *degree of control over collective risk* variable to one: full (single-player game).

Even with only 6 factors left (each with two levels), building a limited, yet useful and balanced, design was no straightforward enterprise. It is worth mentioning that at the end of November 2009 we had some trial experiments where we tested some preliminary ideas and choices. Based on the feedback received then and some further thought, the experimental design was subsequently revised two more times.

Ultimately, we decided to only experiment with *Risk Mode B* (see Appendix C) and to focus on the *malignancy* and *uncertainty* variables. Furthermore, we gave up some of the combinations with the “known function” condition, since – from a “rational” point of view – noise and delay do not matter when the *actual damage* is known. The final design was the one presented in Box 4, with the following advantages:

- every participant only plays 6 games in total that differ significantly from one another; we hypothesize that this will prevent any major learning effects;
- distinguishes between noise and delay, but also studies the joint effects of these two factors;
- still maintains some of the treatments with “known function” and “noise” or “delay”, but mostly for control/exploratory reasons, as no significant effects are expected;
- replication: every treatment is replicated twice, with the exception of Scenarios B’ and C’ which are given less importance (see above);
- ordering effects are controlled for by subdividing each group into two more groups that play the functions in different sequences.

Short Description:

We set up an experiment where each participant has to play 6 similar games, each with different treatment. Each game consists of 10 rounds. Every round, participants receive an income that they have to distribute between a Savings account and a Fund. Investing in the Fund reduces the probability of losing their savings. These probabilities and the information people get depend on the treatment. We are interested in comparing the mean investments for the different treatments, in order to see which factors have a significant effect.

Experimental Design:

	Game 1	Game 2	Game 3	Game 4	Game 5	Game 6
<i>FOR PLAYERS</i>	SCENARIO A		SCENARIO B		SCENARIO A	
	Unknown function			Known function		

<i>GAME INSTRUCTIONS</i>	1 SCENARIO A	2 SCENARIO A	3 Delay // SCENARIO B	4 SCENARIO B	1 SCENARIO A'	2 SCENARIO A'
Session 1 (Group 1)	L	NL	NL	L	L	NL
	NN	NN	NN	NN	NN	NN
	ND	ND	D	D	ND	ND

<i>GAME INSTRUCTIONS</i>	1 SCENARIO A	2 SCENARIO A	5 Noise // SCENARIO C	6 SCENARIO C	1 SCENARIO A'	2 SCENARIO A'
Session 2 (Group 2)	L	NL	NL	L	L	NL
	NN	NN	N	N	NN	NN
	ND	ND	ND	ND	ND	ND

<i>GAME INSTRUCTIONS</i>	3 Delay // SCENARIO B	4 SCENARIO B	7 Noise + Delay // SCEN D	8 SCENARIO C	3 Delay // SCENARIO B'	4 SCENARIO B'
Session 3 (Group 3)	L	NL	NL	L	L	NL
	NN	NN	N	N	NN	NN
	D	D	D	D	D	D

<i>GAME INSTRUCTIONS</i>	5 Noise // SCENARIO C	6 SCENARIO C	7 Noise + Delay // SCEN D	8 SCENARIO C	5 Noise // SCENARIO C'	6 SCENARIO C'
Session 4 (Group 4)	L	NL	NL	L	L	NL
	N	N	N	N	N	N
	ND	ND	D	D	ND	ND

Each of the groups above (20 people/group) was divided in two subgroups (2 x 10 people) that played the functions in different order:

Order A: L, NL, NL, L, L, NL (as in the table above)

Order B: NL, L, L, NL, NL, L.

L = linear function
NL = non-linear function
N = noise
NN = no noise
D = delay
ND = no delay

Factors: 1. known/unknown function
 2. linear/non-linear function
 3. noise/no-noise
 4. delay/no delay

- order of play – possible influences

Dependent variable: Investment in Fund (continuous variable).

In addition, for each participant there is also data on other variables (gathered through a questionnaire). Some of these variables could act as *covariates* – e.g. individual risk attitude – and the data corrected for these influences.

Box 4. Experiment Summary

6.3.2. Participants, Format and Questionnaire

Based on the above mentioned choice of treatments, we conducted the experiment in *four* sessions of *1.5 hours* each. In every session there were *20 participants*, which were randomly selected from a mixed database that is regularly used by other researchers in the field of experimental economics.

From the perspective of the participant, the experiment was structured in three stages and each of the 6 games (called “rounds” in the instructions) could be in one of two possible scenarios. The scenarios defined the informational aspects surrounding the damage level. For instance, in session 1, scenario A referred to a situation where the feedback on the screen could be completely trusted (no noise, no delay), while scenario B explicitly referred to the concept of delay which was linked to the uncertainty aspect:

“The delay is supposed to represent uncertainties from real life. You cannot trust 100% the information you get after every decision, because some of the actions you take now to mitigate climate change may only have a visible effect later” (see Appendix B for further details).

The first stage of every session consisted of two games, one for each scenario, in order to allow the participants to familiarize themselves with the interactive experimental platform and the difference between the two scenarios. This stage did not count for the final earnings. The second stage consisted of four games, two in each scenario. The participants were informed that the conditions of the experiment change between every two games.

Before the third stage, participants received additional information about the shape of the function and the relationship between the number of points in the fund and the actual damage level. Two more games were played in this condition, again one for each scenario.

Last, but not least, while the final individual payment was being counted, a questionnaire was distributed that collected information about the general background of the participants and their experiences with the game.

6.3.3. Game Parameterization

Across all the sessions of this experiment we used the following parameters:

1) *Income and Decision Periods*

We decided that in every game (we call it “round” in our experiment) there would be 10 *decision periods*. Participants were told that every decision period corresponded to one calendar year in real life. Every year each player received an *income* that was fixed at 10 points. They were also informed that this was a discretionary²³ income and that they had to decide how to distribute it between the climate change fund and the personal savings account. The total income for every game was, thus, $I_t = 100$ points.

2) The *linear function* was defined by the following formula:

²³ This means that the annual income received in the game does not represent the full income in any given year, but only what is left after basic expenses are covered (see Appendix B).

$$D_L(F) = 0.95 - 0.009 \cdot F ,$$

where F is the number of points invested in the climate change fund, and $D_L(F)$ represents the damage and returns the probability that all savings ($S=Income-F$) will be lost, when F points have been invested.

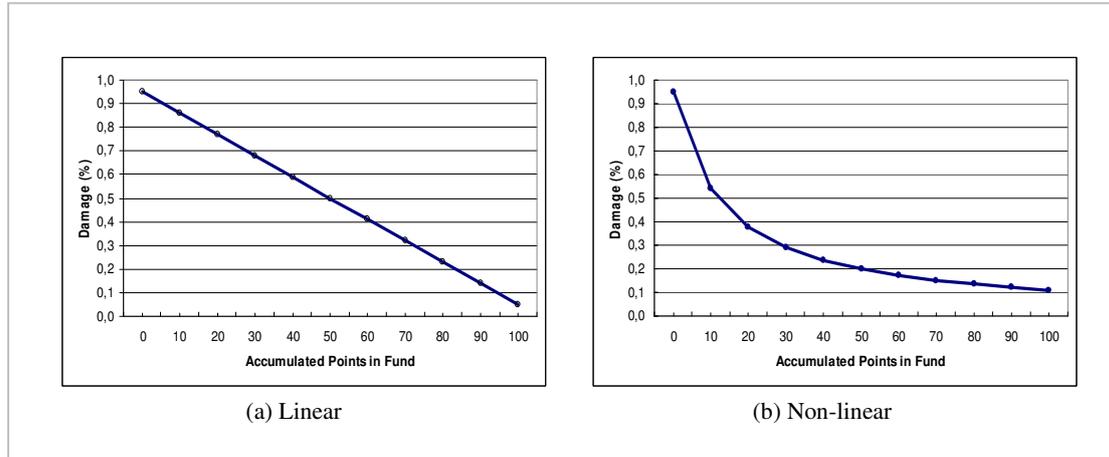


Figure 16. Function Shapes

1) The *nonlinear function* we used was:

$$D_{NL}(F) = \frac{1}{1.0526 + 0.08 \cdot F} ,$$

where – again – F is the number of points invested in the climate change fund and D_{NL} is the probability of loss of all savings.

The coefficients in the function formulas have been selected in such a way that they are comparable on the extremes of the domain. When nothing is invested in the climate change fund, the damage is of 95% - i.e. there is 95% chance that all savings will be lost – and when 90 points out of 100 are invested, the damage for both functions will be between 12-14% (see Figure 16, also Appendix D for larger graphs).

2) *Delay & Noise*

For all rounds that included a treatment with delay and/or noise, we had the following values:

$$delay = 2 \text{ rounds}; \quad noise = 0.2.$$

A delay of 2 rounds means that the players find out about the effect of their decision only two rounds later. For instance, if in the first decision period they invest 5 points in the climate fund, they will only see the corresponding reduction of damage in period 3 (also see Appendix B for further explanations).

The noise parameter sets a fluctuation interval for the random damage levels reported to the players, based on the *actual* damage. In our experiment, this noise is multiplicative and is set at 0.2. This means that when the actual damage is at 75%, for example, the computer will generate a random number that is up to 20% higher or lower than the actual value. This number is then shown on the screen. In all our noise treatments, the players were each time aware of the fact that this is only an *informational* uncertainty and that their final payout will be calculated based on the *actual* value of the damage (without noise).

6.3.4. Outputs: The Data

Experiment

During the experiment, our interactive decision-making platform has saved in a database complete information about the players' decisions at every point in time, the corresponding treatments and the outcomes.

The data available includes, among others: the accumulated investments in the climate fund, the accumulated savings, the probability of loss if the game were to stop at the end of each period (as the output of the corresponding function), as well as the final outcomes – i.e. whether the participant was lucky when the die was thrown and got the full savings or nothing. While it illustrates the potentials of our interactive platform, not all this information will be used in the actual data analysis. What is useful and what is not should be decided for each experiment in question, depending on the risk mode and of the hypotheses that are tested. In our case, more information about what kind of data we used and how will be given in Chapter 8, in the section dedicated to “Data Analysis”.

Questionnaire

The data collected through the questionnaire (see Appendix E) is quite varied and most of it was intended to characterize the sample of participants. Here, we will present the different types of information that we gather and how it is used.

First of all, there are a few questions about the experiment-related strategies and motivations. We report these answers as such and also use them as guidelines for redesigning our hypotheses to be tested.

Secondly, we collect information about some variables that also appear in our conceptual model, as they have been shown to influence risk and solvability perceptions on climate change. Here we refer to: individual risk attitude (*RiskScore* – questions 1-6 in part A), perceived likelihood of climate change (*Likelihood* – question 1 in part B), level of prior information about climate change (*Information* – question 1, part C), personal relevance (*Relevance* – question 2, part C), worry (*Worry* – question 3, part C), government trust (*GovTrust* – question 5, item 1, part C), media trust (*MedTrust* – question 5, item 2, part C), expected effects of climate change (*ExpEff* – items 1-8 in questions 2 and 3, part B). In the statistical model that we use for the experimental data analysis we try to account – where possible – for some of these background variables, introducing them as covariates.

Third, our questionnaire contains items that are not expected to have influenced the decisions in the game, but might be relevant in the context of real-world decisions concerning climate change, such as: the expected monetary consequences of climate change, and the willingness to pay in order to avoid these consequences. We simply report these as characteristics of our sample (see Section 7.2.5.).

Last, but not least, we include a few demographical characteristics to check for the homogeneity of the participants' group and be able to profile the population to which the observed economic behaviour can be generalized. In the next chapter we will further explain how we used this data and what the results tell us about the model we proposed.

7. Experiment Outcomes

7.1. Data Analysis

Before we proceed with the actual results of our study, we address here a few methodological issues that were important in our data analysis. First, regarding the experiment data, we explain here why and how we decided to measure the dependent variable of “individual economic behaviour”.

Secondly, as far as the data from the questionnaire is concerned, we already said earlier (see section 6.3.4) that it served multiple purposes. From the list of items that we included in our questionnaire, some were reported as such – e.g. what percentage of respondents chose which option – while others required some processing before being used in our data analysis. It is also in this subchapter that we discuss more in detail those variables that fall in the last category.

7.1.1. Measures of the Dependent Variable

Earlier when we discussed the general formalization of the dependent variable, we also acknowledged the fact that the exact measures have to be defined with regard to each particular experiment that is carried out using the CLIMEX platform. This is mostly due to the fact that the different risk modes that we introduce translate into different rules of the games, and hence the outcomes are a bit different in every situation.

In the case of our experiment, the participants’ payouts depend highly on chance, as at the end of each game they can get either their full savings or nothing. Consequently, the game score (i.e. the number of points that their monetary payment was based on) would not be an adequate measure for their economic behaviour. Instead, we need to consider some indicators that tell us something about how much of their income were they willing to sacrifice in order to reduce the risk of dangerous climate change. In this sense, there are four possible measures for the dependent variable: a) the total number of points that each participant invested in the fund; b) the expected payout of the chance fork at the end of the game (the chance aspect); c) the deviation from the optimum investment level – i.e. how far away from the “rational” choice he was; d) the final damage level.

Since in our design it doesn’t matter whether players invest money earlier or later in the game, it is also appropriate to mention here that we only analyze the data that describes the situation at the end of period 10 of each round. In other words, we do not evaluate individual behaviour based on the decisions made at every step in the game, but only overall, based on the final results of each game.

a) Level of Investment

For this measure, we use simply the data from the experiment: F is the total number of points invested in the climate change fund, by the end of the 10th period, by each player, for every game, and it is an integer, where $F \in [1,100]$.

b) Expected Payout

The expected payout at the end of each game is calculated according to the expected value formula discussed in Chapter 2, as:

$$EP = p * S ,$$

where S is the number of points accumulated in the savings account by the end of the game and p is the probability to keep the savings.

The points in the savings fund are calculated as the difference between the total income and the accumulated points in the climate change fund:

$$S = I - F ,$$

where $I = 100$ (see Section 6.3.3.).

The probability p to keep the savings is the difference between 1 and the probability of loss, the latter being calculated as an output of the damage function. Thus, as a function of the number of points in the fund, the formula for the expected payout is (see Figure 20 for a graph that shows the expected payout as a function of the total number of points in the climate fund):

$$EP(F) = (1 - D(F)) * (100 - F) .$$

c) Deviation from Optimum

We use the second measure, the deviation from the optimum investment level, as an indication of the level of rationality of the participants. For each game it is possible to calculate an “optimum” level of investment in the climate fund, F_o , such as the corresponding expected payout is maximum. An important mention here is that this optimum can be calculated from the perspective of the experimenter who has full knowledge. However, participants lack information in some treatments, in which case our measure of rationality does not refer to individual rationality, but rather to a collective and absolute type of rationality, relatively to the “best possible option” if there was no uncertainty. Individual rationality is only an issue in those cases where the function is fully known to the participants.

The optimum investment in the fund depends on the damage function and is calculated as the point where the derivative of the payout function is 0 (see Appendix F). In our experiment the optimum investment levels – which render the higher expected payout – are:

- Non-linear function: **24.5** (out of 100).
- Linear function: **47.2** (out of 100).

For each participant and each game, we calculate the normalized deviations from the optimum value:

$$d_f = \frac{F - F_{0f}}{F_{0f}} ,$$

where f is the damage function, $f \in \{linear, nonlinear\}$, F is the total number of points that a player has invested in the climate fund and F_{0f} is the optimum investment corresponding to the function f .

A positive value of d_f will mean that the person has over-invested in the fund (relatively to the “rational” choice), while a negative value will correspond to an under-investment.

d) Final Damage Level

This measure is based on the values saved in the database and is the damage level that each player has reached in the 10th round, based on the accumulated investments in the climate fund.

Which Measures are Most Useful?

Out of the above-mentioned measures, in most parts of our analysis we use the *level of investment* and the *deviation from optimum*. We call these: *CFInv* and *NormDevOpt* respectively.

From the chapter on decision theory we know that probabilities are weighted differently when people are involved in one-time trials, as opposed to repeated trials (see discussion about utilities in Section 2.4.3). Because a die is only thrown once at the end of each game in order to establish whether the participants keep their savings or not, the *expected payout* measure will not capture individual utilities. In the same way, measuring the *final damage level* is only useful as an indication of what is considered a satisfactory (along the lines of Simon 1955) level of risk. In order to avoid the issue of subjective considerations, we decide to concentrate mostly on measures a) and c), thus focusing on the observable investment behaviour and results.

7.1.2. Risk Attitude Profiles

One of the most important variables that we particularly wanted to control for in our study was the individual risk attitude. Since the lottery at the end of each game is a unique event, we needed some way to estimate the utility of the chance fork for every individual and to ensure that we correct our experimental data for the possible variation in risk attitude across groups.

With this objective in mind, we followed the example of Donkers et al. (2001) and included in our questionnaire two types of items: choice questions (CQ) and probability questions (PQ). Based on the participants' answers to these questions we compute *two* risk scores that allow for a ranking of the subjects in terms of risk-aversion. As it will be shown later, we create two risk scores so that we can control one with the other and thus test the reliability of our measure.

Risk Score 1

The choice questions were:

CQ1. Which of the following two options would you choose?

- a. You draw a lottery ticket with a 25% chance of winning 60 euro.
- b. You draw a lottery ticket with a 20% chance of winning 80 euro.

CQ2. Which of the following two options would you choose?

- a. You draw a lottery ticket with an 80% chance of winning 20 euro.
- b. You receive 10 euro for sure.

CQ3. We toss a coin once. Which of the following two options would you choose?

- a. You receive 20 euro with either heads or tails;
- b. You receive 40 euro with heads; but nothing with tails.

For each of the lottery above we computed the expected value, μ , and the standard deviation, σ , i.e. the square root of the variance, of the random variable. The characteristics of each lottery are summarized in the Table 3 below and the calculation details are found in Appendix G. According to Donkers et al. (2001), it is

reasonable to consider the lottery with the higher standard deviation as being the “riskier option”, while the other one is the “safer option”. A person choosing the safer options is considered to be “risk averse”, while a person choosing the riskier options is considered to be “risk seeking”.

Table 3. Lottery Characteristics

Lotteries	Safer Option			Riskier Option		
	EV	Variance	St. Dev.	EV	Variance	St. Dev.
CQ1 60:0.25 vs. 80:0.2	15 (a)	675	25.98	16 (b)	1024	32
CQ2 10:1 vs. 20:0.8	10 (b)	0	0	16 (a)	64	8
CQ3 20:1 vs. 40:0.5	20 (a)	0	0	20 (b)	400	20

Table 4 summarizes which responses to the three questions are risky and which are safe. Also, we present for each question the “rational choice”, i.e. the option with the higher expected value. In the case of CQ3, the expected value is the same for both options, which means that according to rational choice theory subjects should be indifferent between the two.

Table 4. Option Evaluation

	CQ1	CQ2	CQ3
Risky options	b	a	b
Safe options	a	b	a
Rational Choice	b	a	-

The limitation of deriving the risk profile from these questions is that due to differences in utility or other factors, no consistency is to be expected between the answers, thus some people might choose the more risky option in one case and the safer one in a different case. However, for our purposes we do not go into complicated models of estimating risk attitudes, but just want a simple rough measure that allows us to classify individuals into risk-seeking, risk-neutral, or risk-averse. Based on the above, we propose a score where we give 1 point for each question where the safe option was chosen.

If we look at the first two questions, CQ1 and CQ2, we observe that the rational choice is also the “risky” choice. Thus when somebody chooses *b* and *a*, respectively, we do not know for sure that this is because they were being rational (and prefer to maximize the expected value) or because they have a risk-seeking attitude (and don’t mind taking some extra risk for a chance to have a higher payoff). On the other hand, in the case of CQ3, the expected values are the same, thus the choice is clearly an indication of risk attitude and not of rationality. To accommodate for this important difference in our score, we subtract 1 point for choosing the risky option in the case of lottery CQ3. The partial risk score for every option is presented in Table 5.

Table 5. Risk Partial Scores

Lottery	Option (a)	Option (b)
CQ1	1	0
CQ2	0	1
CQ3	1	-1

The total risk score for every individual is obtained by adding the corresponding partial scores for the answers given to each question and is an integer number in the range of [-1, 3]. If our assumptions are correct, a score of 3 should correspond roughly to an individual with a risk-averse attitude, while a score of -1 should correspond to a risk-seeker.

Risk Score 2

The second risk score is calculated based on the answers to the probability questions:

PQ1. Imagine that you just won 10 euro and that you could choose between being paid in cash or receiving a lottery ticket with a certain chance to win a prize of 1000 euro. How high would that chance to win 1000 euro need to be so that you would prefer the lottery ticket to keeping the 10 euro?

PQ2. Imagine that you just won 100 euro and that you could choose between being paid in cash or receiving a lottery ticket with a certain chance to win a prize of 1000 euro. How high would that chance to win 1000 euro need to be so that you would prefer the lottery ticket to keeping the 100 euro?

PQ3. Imagine that you just won 500 euro and that you could choose between being paid in cash or receiving a lottery ticket with a certain chance to win a prize of 1000 euro. How high would that chance to win 1000 euro need to be so that you would prefer the lottery ticket to keeping the 500 euro?

As it can be observed, the subjects were asked to indicate the probability p of winning a certain prize that would be enough for them to give up a sure gain. The underlying assumption is that for the respective probability p the respondent is indifferent between the lottery and the sure gain. If, for each question, a choice of p is made so that the expected values of the two alternatives are equal, the individual is considered to be risk neutral. For each question PQ_i we call this probability p_i^* (see Table 6).

Table 6. Indifference Probabilities

	Lottery	p^*	EV(a)=EV(b)
PQ1	10:1 vs. 1000: p_1	$p_1^* = 0.01$	10
PQ2	100:1 vs. 1000: p_2	$p_2^* = 0.1$	100
PQ3	500:1 vs. 1000: p_3	$p_3^* = 0.5$	500

For every question PQ_i , the higher the probability p_i indicated by a subject, the more attractive the lottery needs to be so that he/she gives up the sure gain, thus the more risk-averse the person is. On the contrary, a low p_i indicates a risk-seeking attitude. By calculating the deviation of p_i from the risk-neutral probability value we can create a rough measure of risk attitude. To account for differences in utility at different probability levels, it makes sense to also normalize these deviations. Consequently, the total risk score is calculated based on the following formula:

$$RS = \sum_{i=1}^3 \frac{p_i - p_i^*}{p_i}$$

Based on these calculations and provided that our assumptions above are true, we create a risk score scale where a higher score indicates a more risk-averse attitude.

Internal Validity

The two measures explicated above should be regarded as approximations of the risk attitudes of the subjects. By creating these scores, our purpose was to make a general risk profile that could be used as a covariate in the statistical model for our experiment. While we were not interested in a fine-tuned measure of risk, we found useful to verify our assumptions by comparing the two risk scores obtained from the different sets of questions. Note that in both cases, a higher score indicates a more risk-averse attitude. Thus, the internal validity of our risk scale could be verified by evaluating the correlation between the two scores. Reliable measures would have to be significantly positively correlated.

The correlation for the data revealed that – indeed – risk score 1 and risk score 2 are significantly related, $r_S = +.270$, $n = 80$, $p < .05$, one tail. The low value of r_S could be explained by the simplicity of our approach in calculating the first risk score. However, despite the weak correlation, the hypothesis test proves that our measures offer a reliable approximation of participants' risk attitudes (also see Figure 17).

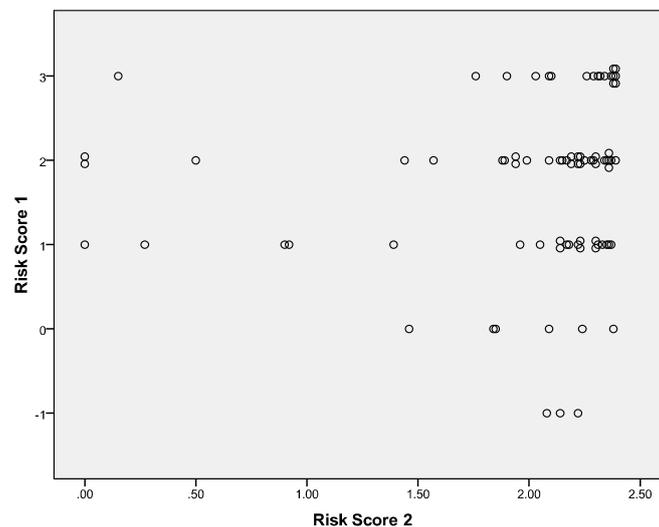


Figure 17. Risk Scores Correlation (risk seeking → risk averse)

External Validity

The only question that still needs to be answered concerns the external validity of our measures – in other words: how do we know that the answers to the questions we used to build the risk scores reflect indeed the real lottery preferences of the participants, especially in the absence of incentives?

First of all, these lottery questions were answered after having participated in the experiment and before receiving the payment for that experiment. Especially in the case of the probability questions where participants are asked to imagine that they had just won a certain amount of money, there is an amount of realism in the context (they had indeed just won some money in the experiment) that could have triggered truthful answers. Although in the questionnaire it was clearly stated that the answers to these questions would have no effect on the earnings from the experiment, the expected payment might have also increased the subjects' effort to duly fill in the questionnaire.

A more powerful argument, however, comes from the literature, where there is evidence that in the case of simple problems, the presence of incentives does not make any difference (see Beattie and Loomes qtd. in Donkers et al. 2001) in the sincerity of the answers.

The analysis above proves that our risk attitude measures are both internally and externally valid and can reliably provide an indication of how risk averse or risk prone any participant in the experiment is. Since they are more or less equivalent, we decided not to aggregate the two measures. Instead, in our analysis of the experimental data, we used the second score, *RiskScore2*, as it has a wider scale and it allows for a better differentiation between risk attitudes.

7.1.3. Other Covariates

Apart from the risk attitude, we also used other covariates in our statistical model, as described in the previous chapter (see Section 6.4.4.). Due to timing constraints in the collection of the data, and because we were expecting that other background variables would play a smaller role in the decision-making processes than risk attitude, we placed less emphasis on controlling the responses in different ways and sometimes we limited the items in the questionnaire to one per variable. This, of course, brings in the danger of unreliable measures. However, the monetary incentives linked to the experiment that we were discussing earlier partially eliminate the risk of incomplete or superficial answers. At the same time, other studies have shown that even with a limited number of questions some useful insights can be obtained. In particular, we were inspired in this endeavour by O'Connor et al. (1999), from which we also borrow the question about likelihood (see below).

Likelihood, Information, Worry, Media Trust and Government Trust

For each of these variables we ask participants to give a score on a 1 to 5 scale, as following:

- *likelihood* that the average annual global temperature will increase by 2°C within the following 50 years: 1 – very unlikely to 5 – very likely;
- how *informed* about climate change the participant thinks he/she is: 1 – not at all informed to 5 – very well informed;
- (personal) *relevance* of the problem of climate change: 1 – very irrelevant, to 5 – very relevant;
- how *worried* the participant is that he/she might be affected by climate change in the future: 1 – not at all worried to 5 – very worried;
- *trust in the government*, measured as level of agreement with the statement “The government will inform me if there is a serious risk that I could be affected by climate change.”: 1 – strongly disagree (corresponding to low government trust) to 5 – strongly agree (high government trust);
- *trust in the media*, measured as level of agreement with the statement “When reporting on climate change issues mass-media is usually neutral.”: 1 – strongly disagree (low trust) to 5 – strongly agree.

These scores are directly used in the statistical model.

Expected Effects

To estimate the perceived expected effects of climate change, we use the answers given to two questions in part B of the questionnaire (Appendix E), each with 8 items. The first question asks participants to rate on a 1 to 5 scale the likelihood of eight different effects of climate change, while the second one asks them to rate on a 1 to 5 scale the magnitude of these effects, assuming that they do happen. Half of the items refer to effects to be incurred by “self”, and half to effects to be incurred by “others”. Based on all these 16 ratings we compute an expected effects score that has a similar structure to that of an expected value:

$$EE = \sum_{i=1}^8 l_i \cdot m_i ,$$

where i denotes one of the 8 effects listed, l is the likelihood of it occurring and m is the magnitude, l and m integers, $l, m \in [1, 5]$.

Such a score is calculated for each participant and controlled for in the experiment data analysis.

7.1.4. Other Indicators: Income Effects

Our questionnaire also included some data about monetary effects of climate change and willingness to pay in order to avoid such effects. Although we do not use this information as input for the experimental data analysis, some explanation about how we quantify the answers to these questions are needed.

First, question 4 of part B of the questionnaire (see Appendix E) asked participants what percentage of their income they think they would have to pay if a certain effect of climate change occurs. The answer options were: below 5%, between 5% and 25%, and above 25%. Similarly, question 5 asked what percentage of their income participants would be willing to invest in a climate change fund in order to avoid certain effects. The same answer options were given. Moreover, the effects that were listed under each of the two questions were the same as those used to calculate the “expected effects” score.

In order to combine these interval measures into one index for the whole group, we used the following approach:

$$I_M(k) = \frac{a \cdot 2.5 + b \cdot 15 + c \cdot 62.5}{a + b + c}, \quad a + b + c = N * k ,$$

where $I_M(k)$ is the index corresponding to k items from the list, a , b and c are the total number of answers in each interval (summed up for the k items), N is the total number of respondents, and the values of 2.5, 15 and 62.5 represent the medians of the intervals [0..5], [5..25] and [25..100] respectively.

The formula above can be used both for the expected monetary effects of climate change and for the willingness to pay to avoid those effects. Also, such an index can be calculated separately for the items concerning the climate change effects that are relevant for “self” and those that are relevant for “others”. In our study, these indexes will only be reported, along with other variables, to characterize the groups that participated in the experiment.

7.2. Results

In this section we reiterate the hypotheses that we started this journey with and present the main results of our experiment.

From the very beginning, we designed our experiment in such a way that the participants would only have to play a limited number of games, all very different from one another, so that learning effects are avoided. Since we were aware that assuming the independence of all trials in the experimental design would simplify our data analysis, we first tested to see that indeed this was the case and that there were no significant differences between the results that could be attributable to the game number. That there was no consistent increasing or decreasing trend in the mean investments is already observable from Figure 18. An analysis of variance on the *CFInv* measure revealed that, indeed, there were no significant learning effects $F(1,111)=.653$, $p=.625$, $\alpha=.05$, $\eta^2=.06$ (see Appendix H, 1). Thus, for the rest of the discussions we can safely consider that the trials are independent.

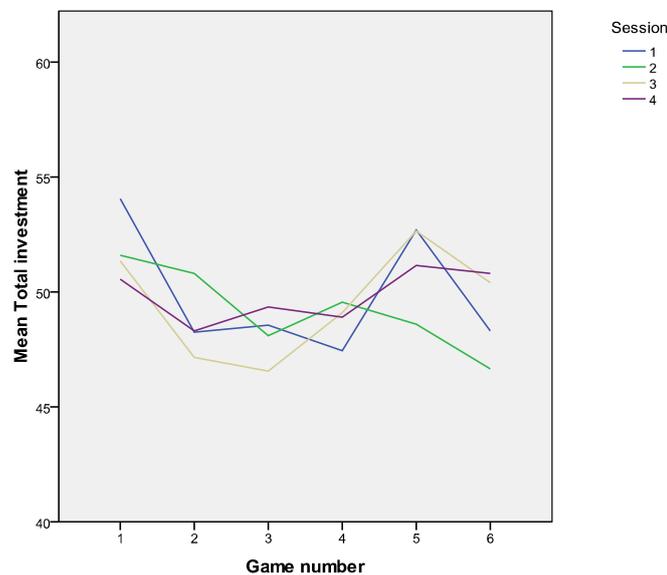


Figure 18. Mean Total Investment in the Fund across the 6 Games, for each Experimental Session

7.2.1. Rationality and Malignancy

One of our hypotheses referred to the fact that people would be more “rational” when the malignancy of climate change is higher. The underlying justification was that the more difficult function shape requires more financial effort in order to reduce the risk significantly and consequently participants will be more concerned with the return on their investment in the climate fund. In plain words: if a lot of money is thought to be at stake, one becomes more calculating and prudent.

A first look at the descriptive statistics in Table 7 reveals that the mean *NormDevOpt* is always higher than 0, which means that overall, participants over-invested in the climate change fund regardless of the malignancy of the function.

Despite the normalization in our rationality measure (see Section 7.1.1) and the fact that an analysis of variance would easily prove the above difference in *NormDevOpt* means significant, it is quite difficult to draw solid conclusions about the effect of malignancy on rationality just by making this group comparison. The

main reason for this is that the two optimum investment levels are not comparable relatively to the total income. While *NormDevOpt* does tell us something about how far a certain investment level is from 47.2 and 24.5 respectively vis-à-vis each of these points, it does not accommodate for the difference in *opportunity* of overinvestment in the two functions. In other words, once a participant has reached the optimum investment level for the non-linear function, he/she has more points left to invest (from the total of 100) than when reaching the optimum investment level for the linear function, an opportunity that might lead to higher over-investment in the former than in the latter case. The same can be argued for an increased opportunity to under-invest in the case of the linear function.

Table 7. Descriptive Statistics: NormDevOpt

Malignancy	Mean	St. Dev.	N	Corresponding Mean CFInv
Linear function	.1977	.31195	240	56.53
Nonlinear function	.7432	.5888	240	42.71

It is for these reasons that we have to keep any analysis of rationality and malignancy at a descriptive level. Nevertheless, interesting observations can be made when we divide the observations of *CFInv* in three categories: “Underinvestment”, “Optimal investment”²⁴ and “Overinvestment” (see Figure 19). While any differences in the frequency of overinvestment and underinvestment between the two functions could be attributed to the above-mentioned opportunity issue, the chart also suggests that the optimal investment level in the climate change fund will be more often reached in the case of the linear than of the non-linear function.

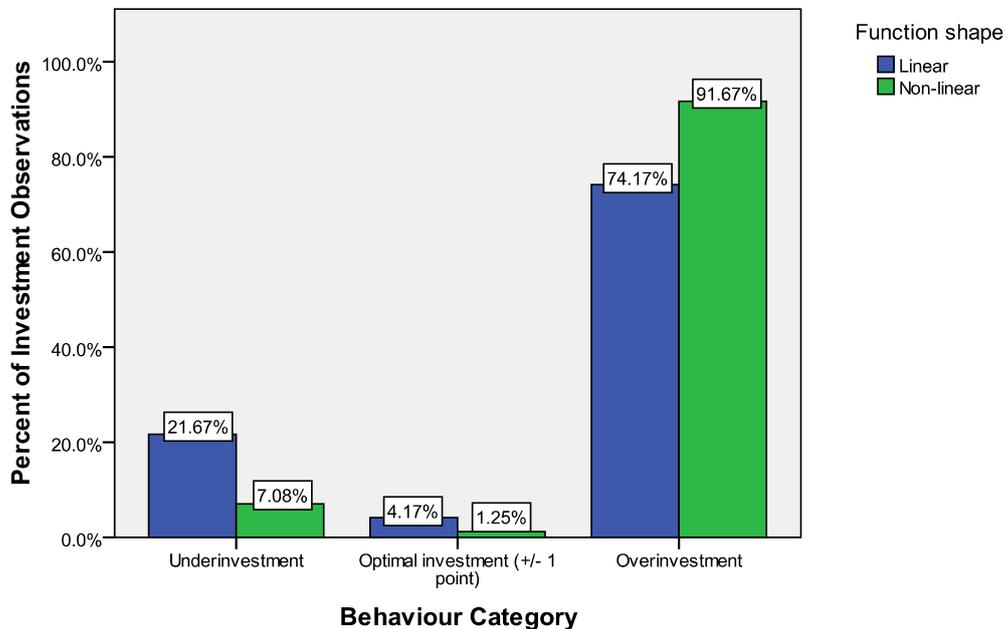


Figure 19. Investment Behaviours and their Relative Frequency

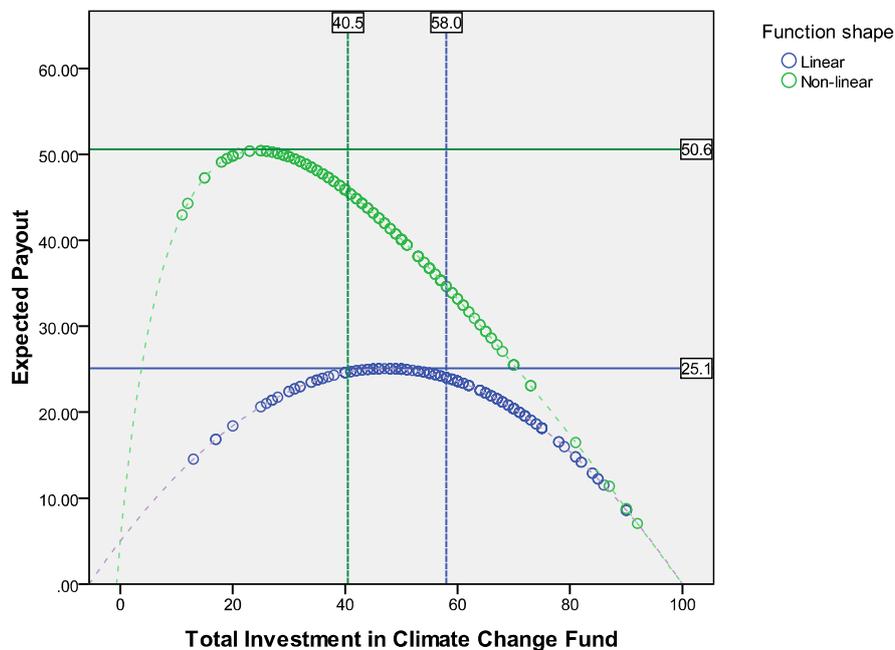
²⁴ For the “Optimal investment” category we include a margin of error of +/- 1 point, due to the fact that the mathematical optima of the functions are calculated with decimals, while investments in the fund can only be integers.

A similar conclusion can be drawn from Figure 20, where we plot the number of points invested in the climate change fund, against the expected payout. Here we can see that there are more points in the vicinity of the maximum expected payout of the linear function than in the vicinity of the maximum expected payout of the non-linear function.

Based on these two observations we can conclude that, in the case of our experiment, our hypothesis is confirmed, as more people played rationally and optimized the expected payout in the games where malignancy was higher. Reasoning from a psychological point of view, one may hypothesize that the linear function (i.e. the situation where the climate problem is malignant) gives players the experience that it is costing a lot of money to do something about it – hence “I become careful and calculating, or even reluctant and fatalistic”. If, on the other hand, the problem is not so difficult to solve, the player notices that throwing only a bit of income at it gives already a significant risk reduction, and consequently his response is “This is easy, I can afford to put in more money and reduce the risk more than I had thought.”

However, since it is not possible to meaningfully compare the groups, nothing can be inferred about the population and the relationship between malignancy and rationality in general.

Further research that will explore this link should ensure that the damage functions are built in such a way that the optimum investment levels are comparable *relatively to the total income*. This could be solved in two ways: either both of the functions are parameterized so that the maximum expected payout is obtained for both of them at the same level of investment in the climate fund (i.e. same F_0 point), or the income scale of the games where malignancy is higher is adjusted so that it compensates for the differences in the optimum CFinv values of the two functions (the two F_0 points of the functions) and is proportionally wider.



The horizontal lines mark the maximum possible expected payouts for each function. The vertical lines indicate the median value of the total investments in the climate fund. Each median line has the colour of its corresponding distribution.

Figure 20. Total Investment vs. Expected Payout

7.2.2. Rationality and Uncertainty

A second hypothesis that we made was that individual economic behaviour would be closer to the rational choice when uncertainties are lower. In the data analysis we use our three-dimensional variable of uncertainty as a single one with seven levels (*uncertainty treatments*). Table 8 gives a summary of the descriptive statistics for the *NormDevOpt* measure across different uncertainty treatments. Because of the reasons outlined in the previous section, we split the cases based on malignancy and compared the groups that had the same function shape.

A one-way analysis of variance reveals that there is no significant difference between the uncertainty treatments, neither in the case of the linear function ($F(6,233)=.1463$, $p=.192$), nor in the case of the nonlinear one ($F(6,233)=.446$, $p=.848$). Moreover, a post-hoc Bonferroni test also doesn't find significant differences between any two groups (see Appendix H, 2). Even when we include other factors in the analysis (see Appendix H, 3), the uncertainty treatment still appears to play no significant role on the deviation from the optimum investment variable. To further explore the effect of uncertainty on the economic investment behaviour of people, we will need to use a more robust statistical model that includes all the variables in our conceptual model. This will have to be done on the *CFInv* measure, because then we will be able to also compare the different malignancy treatments. Consequently, any new conclusions will tell us something about the relationship between uncertainty and the level of investment, not rationality.

For the moment we accept that the level of uncertainty has no effect on the rationality of economic behaviour. Although we were expecting that most people will be optimising their investment when they *know* how the damage functions work and feedback information is easily *available* and *reliable*, it turns out that this is not the case. In a sense such a finding is not all that surprising, given that behaviour decision theory has consistently shown that rationality assumptions are often violated in practical decision-making.

Table 8. Descriptive Statistics for each Uncertainty Treatment: NormDevOpt

<i>Uncertainty Treatment</i>	<i>Linear Function</i>			<i>Nonlinear function</i>		
	<i>Mean</i>	<i>St. Dev</i>	<i>N</i>	<i>Mean</i>	<i>St. Dev</i>	<i>N</i>
Baseline (K=Known) Known function, No noise, No delay (000)	.2256	.29686	40	.6439	.53442	40
Baseline (U=Unknown) Unknown function, No noise, No delay (100)	.2505	.30041	40	.7684	.54631	40
Known						
Delay (K) No noise (010)	.2913	.24249	20	.7184	.57843	20
Noise (K) No delay (001)	.2617	.29409	20	.7306	.47809	20
Unknown						
Delay (U) No noise (110)	.1811	.30517	40	.6939	.65749	40
Noise (U) No delay (101)	.1430	.36551	40	.8082	.62575	40
Delay + Noise (U) (111)	.1091	.31434	40	.8204	.64667	40

However, in the case of this particular experiment, where the decisions made wouldn't have had any impact on other people nor on the external environment, considerations of social utility (see Camerer 2004 and Section 3.3.2) shouldn't have played a role in the rationality of play. A more plausible explanation for the consistent overinvestment, even in the cases where the function was known, would be related to framing effects.

Tversky and Kahneman (1981), among others, have shown that the ideas that decision-makers have about their actions and the associated outcomes influence greatly the choice made. Thus it is very probable that the fact that the decision problem was framed as climate change has led to over-investments in the fund. Even though participants knew that their decisions in the game would not be related in any way to real-life climate change involvement, they might have felt morally compelled to invest more in the climate change fund than they would have in the case of a de-contextualized game.

This is not bad news, but on the contrary, as we particularly wanted to know how people would behave *in the context of climate change*. From the questionnaire we know that our group of participants reported a high likelihood score that an increase of the global annual temperature beyond 2°C would happen in the next 50 years, with an average of 4.10 points out of 5. We could thus hypothesize that uncertainty does not matter that much for investment behaviour *once* people have already accepted the idea that climate change is an ongoing phenomenon that might affect them in the near future.

To verify this, future research could try to study behaviour under the same uncertainty treatments, but also include a control group in which the decision-problem is not contextualized. Also, an idea would be to test subjects' beliefs and attitudes toward climate change before the experiment and randomly distribute them into a group of "believers" and one of "sceptics".

7.2.3. Rationality and Order

Although we had initially not made any hypothesis about ordering effects on the economic investment behaviour, our experimental design controlled for such an effect of the malignancy of the function, by dividing each group of people in two subgroups. Thus, in every session, 10 people played games in which the malignancy of the function alternated in the following pattern: L, NL, NL, L, L, NL (order A) while 10 people played NL, L, L, NL, NL, L (order B). When creating these sequences we took into account the fact that every uncertainty treatment has two games with different functions. By rotating the order of the functions in the second uncertainty treatment (games 3 and 4) we wanted to make sure that we don't create a recognizable pattern of malignancy either, hence avoiding repetition effects.

The surprising finding was that order had an effect on the individual economic behaviour. Our analysis of variance on *NormDevOpt* (see Appendix H, 3) showed a weak, but significant effect of order, $F(1,452)=9.240$, $p=.003$, $\eta^2=.020$ at $\alpha=.05$. Because the model also included malignancy as an independent variable and we have already established that we should not compare *NormDevOpt* across different malignancy treatments, we also ran two analyses of variance, one for each function shape (see Appendix H, 4). This gave us a clearer picture of the influence of order on rationality, by revealing that order only had an effect on rationality in the treatments with the non-linear function, $F(1,226)=8.659$, $p=.004$, $\eta^2=.037$ at $\alpha=.05$. The mean

deviation from the optimum investment under the non-linear treatment was significantly higher for people in subgroup B – $M=.857$ – than for people in subgroup A – $M=.624$ (also see Figure 21).

In other words, when playing a game with a non-linear function, those who had started the game also with a non-linear function were less rational than those who had started with the more difficult, linear function. Although not significant statistically, a similar pattern is observable also when playing a game with a linear function.

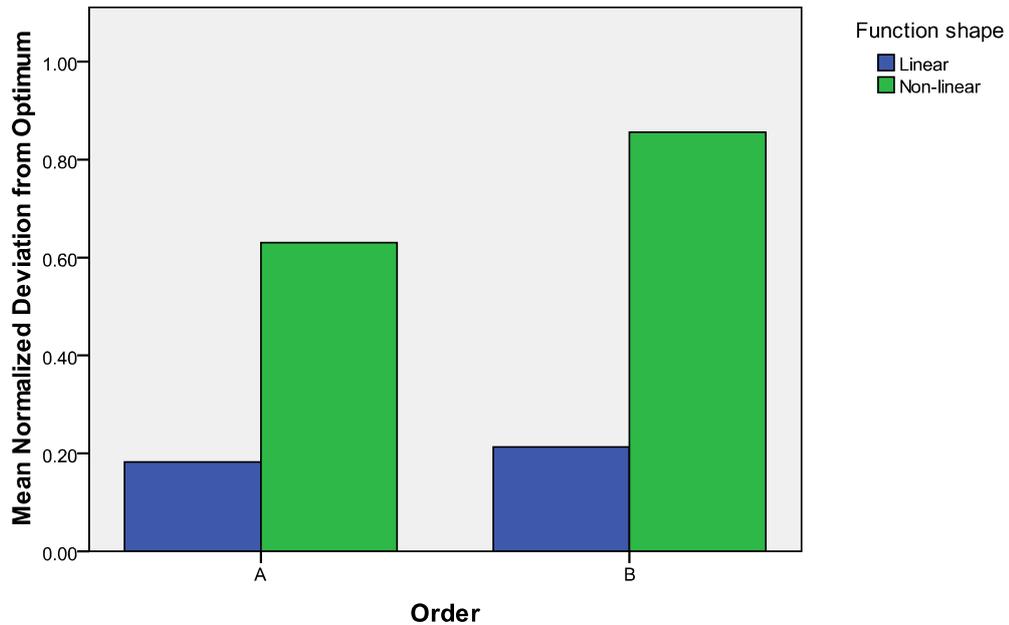


Figure 21. Mean Normalized Deviation from Optimum for different Order and Malignancy Treatments

This phenomenon might be linked to the solvability perception variable in our model. More specifically, when confronted first with an easier problem, people might enter a state of mind where they feel they can cope better with any type of problem. Such a hypothesis is suggested by a participant from the A subgroup who expressed his frustration when confronted with the difficult linear function, especially since they had first played some trial games where the level of malignancy was comparable to the one of the non-linear function – and considerably easier:

“in the first round I was slightly put off because the trial rounds had been so easy (just a few points in CC lowered the probability a lot), while that first "real" round was decidedly harder (it cost a lot of points to lower the probability). I believe that was the greatest chance I took (+/-50%).”

The effect of previous experience is also brought into discussion by a player in subgroup B:

“I think that the morale after each round also played a role in establishing my decisions for the following rounds. I don't know if this was intended, but it was definitely there.”

If this previous game experience has had an effect indeed, then we should also be able to observe an order effect for *CFInv*, with people playing in order B investing on average more than those in order A, regardless of the optimum investment value.

Before we shape a more robust explanation, we have to thus analyze the level of investment in the following section.

Nevertheless, a first recommendation for further research that we can already make would be to include in the conceptual model a variable relating to experience. This could be formalized either as “recent experience” where we look each time at the malignancy in the game before, or “first-time experience” where – like in the case of order here – we focus on the malignancy in the first game. The latter formalization could also be linked to some already existent proof in the literature about a so-called *status quo bias*. Kahneman et al. (2000) discuss this effect from the perspective of loss aversion, where people have a strong tendency not to change their once-tested strategy (or current situation), as they weight the possible disadvantages of doing so significantly more than the possible gains.

7.2.4. Effects on Investment Level

So far we have discussed the effects of different variables on the rationality of the decisions made by participants. However, due to the characteristics of the *NormDevOpt* measure we were not able to make reliable comparisons across different treatments of malignancy. Now, we use the investment level as our dependent variable (*CFInv*) and we can run a more integrated statistical model.

Since we also wanted to include some of the background variables in our model as covariates in our analysis, we first tested the homogeneity of regression assumption that is needed for an ANCOVA (see Appendix H, 5). Based on this analysis, we eliminated the background variables that interacted significantly with at least one of the independent variables, in this case *ExpEff* and *Worry*. Then we ran the analysis of variance test, but because some of the covariates had no significant effect, we discarded those results and decided to further eliminate from our analysis *Information*, *Likelihood* and *Relevance*. If such measures are to be included in a statistical model in the future, more attention should be paid to them also in the questionnaire. In our case, it seems that a simple item for each of these variables was not enough to get a reliable measure. On the other hand, the *RiskScore* covariate, that we most thoroughly measured, was a very useful one. Lastly, while *GovTrust* and *MedTrust* appeared significant in our model, their effects should be interpreted with caution, since again we only used in the questionnaire one item for each.

The final model on *CFInv* gave results that were similar to those related to rationality. The uncertainty variable did not have a significant effect on the level of investment in the climate change fund, $F(6,449)=.558$, $p=.764$, $\alpha=.05$. Also, not even when we compared the uncertainty treatments two by two, we couldn't find any significant difference.

The malignancy variable, however, resulted in significantly higher investments in the climate change fund for the linear than for the non-linear function, $F(1, 449)=.114.861$, $p<.001$, $\eta^2= .204$, $\alpha=.05$. This tells us, first of all, that participants did not anchor on a certain value of the level of investment, regardless of the game they were playing.

Given also the observed difference in rationality based on the malignancy of the function, we wondered if participants had instead anchored on a certain level of damage, or better yet on a certain expected payout, across the different malignancy treatments. Table 9 suggests that this was not the case, as there were large differences

in the means obtained for the two different malignancy levels, across all the measures of the dependent variable.

Table 9. Descriptive Statistics for Different Measures of the Dependent Variable and Functions

<i>Dependent Variable Measure</i>	<i>Linear Function</i>			<i>Nonlinear function</i>		
	<i>Mean</i>	<i>St. Dev</i>	<i>N</i>	<i>Mean</i>	<i>St. Dev</i>	<i>N</i>
Rationality (<i>NormDevOpt</i>)	.1977	.31195	240	.7432	.5888	240
Total Investment (<i>CFInv</i>)	56.53	14.724	240	42.71	14.426	240
Final Damage Level	44.21	13.266	240	24.00	6.606	240
Expected Payout	22.34	3.300	240	42.69	7.938	240

Coming back to our integrated statistical model, another significant effect on the level of investment in the climate change fund was that of order, $F(1,449)=9.791$, $p=.002$, $\eta^2=.021$, $\alpha=.05$. This confirms the idea that we discussed earlier that previous experience does play a role in current decisions, even if indirectly. This is not about actual learning from different uncertainty treatments (we had shown already that there was no effect of the game, see the beginning of this section), but about a solvability perception depending on the level of difficulty of the problem at hand, where some emotional component might play a role. To explore this phenomenon a bit further, we created a new variable “Malignancy in the previous round”, with three possible values: 0 for no experience (first round), 1 for the games in which the recent experience had been with a linear function, and 2 for the games where a non-linear function had been played in the round before. Without making any significance estimates, it is interesting to observe the pattern of total investments in Figure 22.

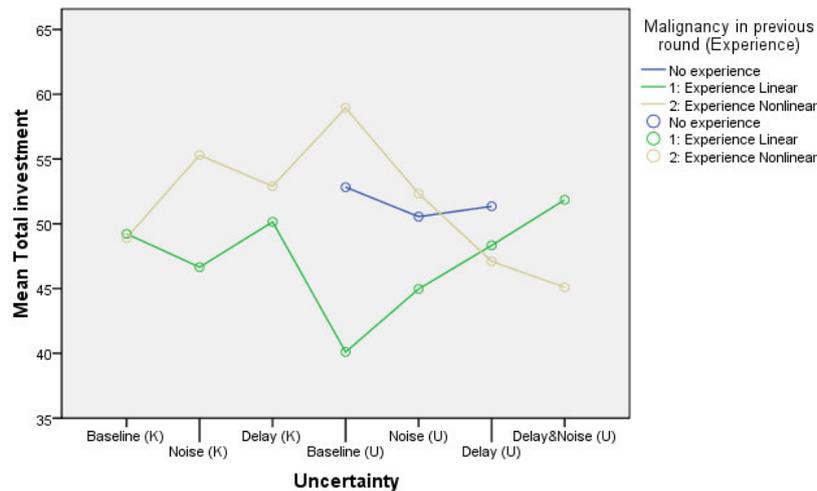


Figure 22. Malignancy in Previous Round and Uncertainty (no significance estimates)

Along with our other findings, this graph suggests once more that any further study about decision-making, especially in the context of climate change, should pay special attention to the variables that could influence the solvability perception.

As far as the covariates in our model are concerned, as we would have expected, the risk score is an important predictor of the level of investment in the climate change fund (for coefficients, see Appendix H, 6, Parameters). The more risk averse an individual, the more he invested in the climate change fund. On the other hand, the level of trust in the government and media seems to affect investment negatively (i.e. the more people trust the government or the media, the less they invest), but as

mentioned before, these results should be interpreted with care, because of the simple way in which the *GovTrust* and *MedTrust* variables were measured.

7.2.5. Characteristics of the Participants

Before we put all these findings together and conclude our report, we would like to make here a few more remarks about the characteristics of the participants.

First of all, as it is often the case in this kind of experiments, our samples consisted mainly of Dutch university students. This brings in some obvious limitations with regard to generalizing these findings to a more diverse population. Nevertheless, since this type of research that uses economic experiments under uncertainty in the context of a particular decision situation is very new, our study needs to be interpreted as an exploratory one which will provide the basis for further, more refined experiments. A surprising thing is that our questionnaire revealed that about 30% of the participants were economics students, yet the group deviates considerably from the “rational” choices. In a more diverse group, though, we could expect that the deviation would be even greater, but this is something to investigate in the future.

Regarding other demographic variables, 58% of the participants were female and 42% male, the average age was 22, and 90% had an income level below 1500€/month. Additional information about other variables in the questionnaire is summarized in Table 10.

Table 10. Questionnaire Results

<i>Variable</i>	<i>Mean</i>	<i>St. Dev.</i>	<i>N</i>
RiskScore2	41.67	29.08	80
Likelihood (1 to 5 scale)	4.10	1.03	80
Expected Effects Score	39.89	15.63	80
Expected Effects Score	64.16	16.98	80
Expected Monetary Consequences	96.2	N/A	80
Expected Monetary Consequences	44.38	N/A	80
Willingness to Pay to Avoid Consequences	80.73	N/A	80
Willingness to Pay to Avoid Consequences	50.18	N/A	80
Information (1 to 5 scale)	3.61	1	80
Relevance (1 to 5 scale)	3.76	1.16	80
Government Trust (1 to 5 scale)	3.48	1.19	80
Media Trust (1 to 5 scale)	2.25	1.10	80

The calculation of the income effects (“Expected Monetary Consequences” and “Willingness to Pay to Avoid Consequences”) is described in Section 7.1.4. We provide these scores here only as an indication of how people expect the consequences of climate change to translate in monetary terms, both relatively to their situation and that of other people elsewhere in the world.

8. Conclusions and Recommendations

Solving the climate change issue is by no means an easy task – this is an idea that we had in our mind when we decided to walk the research path sketched in this report. Our purpose was to study individual economic behaviour under risk and uncertainty, in a setting that is specifically contextualized as climate change. As stated in the introduction, by carrying out this research we aimed at *further clarifying the policy problem of climate change*.

The extent to which we managed to achieve this objective can only be judged relatively to the broader context of this study. More specifically, it must be stated here that this study was conducted as part of the Global System Dynamics (GSD) and Policies project, funded by the Future & Emerging Technologies division of the European Commission²⁵. Since applying the lessons of experimental economics on such a specific topic such as climate change was a novel pursuit in itself, being part of a larger project meant that we could safely choose to be explorative in our approach. In other words, we could afford not to have only narrow and clear-cut hypotheses and conclusions, but also to take the time to develop a framework that could be used in further research. This is not to say, however, that more to-the-point observations were absent from our study; on the contrary, we have already stated some of the most important ones in Section 7.2.

After reviewing an interdisciplinary body of literature we were able to build a conceptual model that could be particularly suited to the study of climate change related decisions. Moreover, the development of the complementary interactive decision-making platform, CLIMEX, could be well regarded as another success of this research.

Due to time constraints, many of the links provided in our model could not be explored. It remains for future similar endeavours to try to verify those links and improve the model when adequate. The limited laboratory experiments that we conducted had already indicated where such improvements could be made, one of the most important points being that of further specifying the factors that influence the solvability perception of decision-makers (see Sections 7.2.3 and 7.2.5 for details about the order effects).

While we did not have the time and resources to also conduct some experiments in a multiple-player condition, this possibility is now open for the very near future. In a recent GSD workshop²⁶ we had a chance to already try out the new feature of CLIMEX that allows the experimenter to design a game in which more participants will invest in a common climate change fund, thus creating a public good dilemma. Currently, discussions exist about the possibility of opening the platform so that such experiments could be run through Internet on a wider scale.

One of our ambitions at the beginning of this study was to also try to make some recommendations for policy makers engaged in the formulation of climate change related policies. At this point in time and based on our work so far, one of the

²⁵ www.globalsystemdynamics.eu

²⁶ GSD Workshop 'Elementary models for a sustainable economy', held at Utrecht University, 21-14 January 2010.

most interesting things we can conclude is that it appears that the uncertainties associated with climate change actually play no role in economic decision-making. Our research has showed that, when it comes to behaviour, what people *believe*, and how they *perceive* the current situation is more important than rational considerations, even when uncertainty is not an issue.

From the perspective of a possible climate change regime, or even of local level policies, the fact that people are not very rational in the face of risk could actually be useful for finding the needed political support. If they can be convinced that climate change does expose us to a risk, then individuals will be willing to take at least some action to reduce it. Of course, the question remains whether the level of involvement will be sufficient to actually avoid irreversible and possibly catastrophic changes in the climate. In this sense, we think governments – in agreement with scientists – should take the lead and be the ones to define and impose what is a “safe” level of risk, as it happens now with the precautionary principle of the EU and the commitment to keeping the global annual average temperature below 2°C.

The good news is that financial support for such a policy could be found, although often times it will probably be accompanied by complaints. While it is true that variables such as *solvability and risk perception* do play a major role when it comes to *how much* people will pay in order to avoid a (known, or unknown) risk, they will still continue to reach for what each of them considers to be a comfortable level of risk, regardless of how costly that might be. Our study has shown that people are highly adaptive to the present situation: in our results (Figure 18, for instance) we saw that generally the investment levels have always been suited and comparable to the malignancy of the function, even when people might have been “put off” – in the words of a participant – by the difficulty of the situation. What *is* important is that people have an idea about this malignancy – i.e. they are continuously fed even rough estimates of the risks they run if they don’t act. The precision of those estimates is way less important.

For policy makers this translates into a recommendation that they keep the climate change debate alive, no matter what. Even when different interests might and will try to discredit the scientific facts underlying climate change, those who want a policy resolution to this major environmental problem should continue to communicate the message of risk exposure. It doesn’t matter if there are also sceptics in this world, show must go on...

References

- Ahn, T. K.; Ostrom, Elinor; Schmidt, David; Shupp, Robert and James Walker (2001) "Cooperation in PD Games: Fear, Greed and History of Play" *Public Choice* 106: 137-155.
- Ambec, Stefan and Lars Ehlers (2008). "Cooperation and equity in the river-sharing problem" in: Dinar, Ariel; Albiac, José and Joaquín Sánchez-Soriano (eds.) *Game Theory and Policymaking in Natural Resources and the Environment*. Routledge: New York, pp. 112-131.
- Andreoni, James and John H. Miller (1993). "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence". *The Economic Journal* 103(418), pp. 570-585.
- Becker, Selwyn and Fred Brownson (1964). "What price ambiguity? Or the role of ambiguity in decision-making." *Journal of Political Economy* 72, pp. 62-73.
- Binmore, Ken (2007). *Playing for Real – A Text on Game Theory*. Oxford University Press: Oxford, New York.
- Blowers, Andrew (1998). "Power, participation and partnership", in: P. Glasbergen (ed.) *Cooperative Environmental Governance. Public-private agreements as a policy strategy*, Kluwer Academic Publishers, pp. 21-42.
- Brundtland Commission, United Nations (1987). "Our Common Future". Report of the World Commission on Environment and Development, General Assembly Resolution 42/187, 11 December 1987. Available at: <http://www.un-documents.net/wced-ocf.htm>. Accessed November 2009.
- Byron, B. Michael Jr. (1996). *Rationality and the Paradoxes of Decision Theory: A critique of Rational Choice Views*. UMI: Ann Arbor.
- Cadernás, Juan Camilo & Elinor Ostrom (2004). "What do people bring into the game? Experiments in the field about cooperation in the commons". *Agricultural Systems* 82(3), pp. 307-326.
- Camerer, Colin F. (1995). "Individual Decision Making" in: Kagel, John H. and Alvin E. Roth (eds). *The Handbook of Experimental Economics*, pp. 587-673.
- Camerer, Colin F. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press: Princeton, NJ.
- Camerer, Colin (2004). "Behavioural Game Theory: Predicting Human Behavior in Strategic Situations" in: Camerer, F. Colin; Loewenstein, George and Matthew Rabin (eds.). *Advances in Behavioral Economics*. Princeton University Press: New Jersey, pp.374-392.
- Cameron, Trudy Ann (2005). "Updating Subjective Risks in the Presence of Conflicting Information: An Application to Climate Change". *The Journal of Risk and Uncertainty* 30(1), pp. 63-97.
- Cardenás, Juan-Camilo and Elinor Ostrom (2004). "What do people bring into the game? Experiments in the field about cooperation in the commons". *Agricultural systems* 82(3), pp. 307-326.

- Carter, Michael W. and Camille C. Price (2001). *Operations Research: A Practical Introduction*. CRC Press LLC: Boca Raton, Florida.
- Chakravarty, Sujoy and Jaideep Roy (2008). “Recursive Expected Utility and the Separation of Attitudes towards Risk and Ambiguity: an Experimental Study”. *Theory Dec.* 66, pp. 199-228.
- de Vries, Bert J.M. ; Petersen, Arthur C. (2009). “Conceptualizing sustainable development: An assessment methodology connecting values, knowledge, worldviews and scenarios”. *Ecological Economics* 68, pp. 1006-1019.
- Donkers, Bas; Melenberg, Bertrand and Arthur van Soest (2001). “Estimating Risk Attitudes using Lotteries: A Large Sample Approach”. *The Journal of Risk and Uncertainty* 22(2), pp. 165-195.
- Eek, Daniel and Tommy Gärling (2008). “A New Look at the Theory of Social Value Orientations: Prosocials Neither Maximize Joint Outcome nor Minimize Outcome Differences but Prefer Equal Outcomes” in: Biel, Anders; Eek, Daniel; Gärling, Tommy and Mathias Gustafsson (eds). *New Issues and Paradigms in Research on Social Dilemmas*. Springer: New York, pp. 10-26.
- Faysse, Nicolas (2005). “Coping with the Tragedy of the Commons: Game Structure and Design of Rules.” *Journal of Economic Surveys* 19(2), pp. 239-255.
- Fehr, Ernst and Klaus M. Schmidt (1999). “A theory of fairness, competition and cooperation”. *The Quarterly Journal of Economics: August*; pp. 817-868.
- Finus, Michael (2001). *Game Theory and International Environmental Cooperation*. Edward Elgar: Cheltenham, UK; Northampton, MA.
- Gärdenfors, Peter and Nils-Eric Sahlin (1988). “Introduction: Bayesian decision theory – foundations and problems” in Gärdenfors, Peter and Nils-Eric Sahlin (eds.) *Decision, Probability and Utility: Selected Readings*. Cambridge University Press: Cambridge.
- Hacking, Ian (2006). *An Introduction to Probability and Inductive Logic*. Cambridge University Press: Cambridge.
- Hardin, Garrett (1968). “The tragedy of the commons”. *Science* 162, pp. 1243-1248.
- Hargreaves Heap, Shaun P. and Yanis Varoufakis (1995). *Game Theory: A Critical Introduction*. Routledge: New York.
- IPCC (2007a). Fourth Assessment Report: Climate Change 2007 (AR4). Available online at: http://www.ipcc.ch/publications_and_data/publications_and_data_reports.htm #1. Accessed August 2009.
- IPCC (2007b). AR4, Working Group I, Summary for Policymakers (SPM). Available online at: <http://www.ipcc.ch/pdf/assessment-report/ar4/wg1/ar4-wg1-spm.pdf>. Accessed August 2009.
- Janssen, Marco and Bert de Vries (1998). “The Battle of Perspectives: a multi-agent model with adaptive responses to climate change”. *Ecological Economics* 26, pp. 43-65.
- Kahneman, Daniel and Amos Tversky (2000). “Prospect Theory: An Analysis of Decision-Making under Risk” in: Kahneman, Daniel and Amos Tversky (eds.). *Choices, Values and Frames*. Cambridge University Press: Cambridge, pp. 17-43.

- Kahneman, Daniel; Knetsch, Jack L. and Richard H. Thaler (2000). "Anomalies: The Endowment Effect, Loss Aversion and Status Quo Bias" in: Kahneman, Daniel and Amos Tversky (eds.). *Choices, Values and Frames*. Cambridge University Press: Cambridge, pp. 159-170.
- Kazemi, Ali and Daniel Eek (2008). "Promoting Cooperation in Social Dilemmas via Fairness Norms and Group Goals" in: Biel, Anders; Eek, Daniel; Gärling, Tommy and Mathias Gustafsson (eds). *New Issues and Paradigms in Research on Social Dilemmas*. Springer: New York, pp. 72-92.
- Knez, Marc and Colin Camerer (2000). "Increasing Cooperation in Prisoner's Dilemmas by Establishing a Precedent of Efficiency in Coordination Games". *Organizational Behavior and Human Decision Processes* 82(2), pp. 194-216.
- Leiserowitz, Anthony (2006). "Climate change risk perception and policy preferences: The role of affect, imagery and values". *Climatic Change* 77(1-2), pp. 45-72.
- Liu Hsin-Hsien and Andrew M. Colman (2009). "Ambiguity aversion in the long run: Repeated decisions under risk and uncertainty". *Journal of Economic Psychology* 30, pp. 277-284.
- Milinski, Manfred; Semmann, Dirk; Krambeck, Hans-Jürgen and Jochem Marotzke (2006). "Stabilizing the Earth's climate is not a losing game: Supporting evidence from public goods experiments". *PNAS* 103(11), pp. 3994-3998.
- Milinski, Manfred; Sommerfeld, Ralf D.; Krambeck, Hans-Jürgen, Reed, Floyd A. and Jochem Marotzke (2008). "The collective-risk social dilemma and the prevention of simulated dangerous climate change". *PNAS* 105(7), pp. 2291-2294.
- Morris, Christopher W. and Arthur Ripstein (2001). "Practical Reason and Preference" in Morris, Christopher W. and Arthur Ripstein (eds.) *Practical Rationality and Preference: Essays for David Gauthier*. Cambridge University Press: Cambridge.
- O'Connor, Robert E.; Bord, Richard J. and Ann Fisher (1999). "Risk Perceptions, General Environmental Beliefs and Willingness to Address Climate Change". *Risk Analysis* 19(3), pp. 461-471.
- Okada, Akira (2008). "The second-order dilemma of public goods and capital accumulation". *Public Choice* 135, pp. 165-182.
- Olivier, Pamela (1980). "Rewards and Punishments as Selective Incentives for Collective Action: Theoretical Investigations". *The American Journal of Sociology* 85(6), pp. 1356-1375.
- Ostrom, E., Gardner, R. and J. Walker. (1994). *Rules, Games and Common-pool Resources*. University of Michigan Press, Ann Arbor.
- Ostrom, Elinor; Burger, Joanna; Field, Christopher B.; Norgaard, Richard B. and David Policansky (1999). "Revisiting the Commons: Local Lessons, Global Challenges". *Science* 284(5412), pp. 278-282.
- Ostrom, Elinor (2006). "The value-added of laboratory experiments for the study of institutions and common-pool resources". *Journal of Economic Behavior & Organization* 61: 149-163.

- Ostrom, Elinor (2007). "A diagnostic approach for going beyond panaceas". *PNAS* 104(39), pp.15181-15187.
- Peterson, Martin (2008). *Non-Bayesian Decision Theory: Beliefs and Desires as Reason for Action*. Springer. Electronic version available (by subscription) online at: <http://www.springerlink.com/content/978-1-4020-8698-4>. Accessed June 15th, 2009.
- Pittock, A. Barrie (2009). *Climate change: the science, impacts and solutions*. 2nd edition. CSIRO Publishing: Australia.
- Pulford, Briony D. and Andrew M. Colman (2006). "Ambiguous Games: Evidence for strategic ambiguity aversion". *The Quarterly Journal of Experimental Psychology* 60(8), pp. 1083-1100.
- Pulford, Briony D. and Andrew M. Colman (2008). "Size Doesn't Really Matter: Ambiguity Aversion in Ellsberg Urns with Few Balls". *Experimental Psychology* 55(1), pp. 31-37.
- Raaijmakers, Ruud; Krywkow, Jörg and Anne van der Veen (2008). "Flood risk perceptions and spatial multi-criteria analysis: an exploratory research for hazard mitigation". *Nat. Hazards* 46, pp. 307-322.
- Richardson, Katherine; Steffen, Will; Schellnhuber, Hans Joachim; Alcamo, Joseph; Barker, Terry; Kammen, Daniel M.; Leemans, Rik; Liverman, Diana; Munasinghe, Mohan; Osman-Elasha, Balgis; Stern, Nicholas; Wæver, Ole. (2009). "Synthesis Report" from the International Scientific Congress "Climate Change – Global Risks, Challenges & Decisions", Copenhagen 2009, 10-12 March. Available online at: <http://climatecongress.ku.dk/pdf/synthesisreport/>. Accessed August 2009.
- Russell, Bertrand (1912). *The Problems of Philosophy*. Henry Holt and Company: New York; Williams and Norgate: London.
- Sally, David (1995). "Conversation and cooperation in social dilemmas: A meta-analysis of experiments 1958-1992." *Rationality and Society* 7(1), pp. 58-92.
- Savage, Leonard J. (1972). *The Foundations of Statistics*. Dover Publications Inc.: New York.
- Seip, Knut Lehre and Fred Wenstøp (2006). *A Primer on Environmental Decision-Making: An Integrative Quantitative Approach*. Springer: Dordrecht, The Netherlands.
- Simon, Herbert A. (1955). "A Behavioral Model of Rational Choice". *Quarterly Journal of Economics* 69(1), pp. 99-118.
- Steg, Linda and Charles Vlek (2009). "Social Science and Environmental Behaviour" (chapter 9) in: Boersema, Jan J. and Lucas Reijnders (eds.) *Principles of Environmental Sciences*, Springer Science.
- Sterman, John D. and Linda Booth Sweeney (2007). "Understanding public complacency about climate change: adults' models of climate change violate conservation of matter". *Climatic Change* 80(3-4), pp. 213-238.
- Tversky, Amos and Craig R. Fox (2000). "Weighing Risk and Uncertainty" in: Kahneman, Daniel and Amos Tversky (eds.). *Choices, Values and Frames*. Cambridge University Press: Cambridge, pp. 93-117.

- Tversky, Amos and Daniel Kahneman (1981). "The Framing of Decisions and the Psychology of Choice". *Science* 211 (1181), pp. 453-458.
- Viscusi, W. Kip and Richard J. Zeckhauser (2006). "The perception and valuation of the risks of climate change: a rational and behavioural blend", *Climatic Change* 77, pp. 151-177.
- Vlek, Charles and Keren, Gideon (1992). "Behavioural decision theory and environmental risk management: Assessment and resolution of four 'survival' dilemmas". *Acta Psychologica* 80, pp. 249-278.
- Weirich, Paul (2004). *Realistic Decision Theory: Rules for nonideal agents in nonideal circumstances*. Oxford University Press: New York.
- Winkler, Robert L. (1972). *Introduction to Bayesian Inference and Decision*. Holt, Rinehart and Winston, Inc.: New York.
- Von Winterfeldt, Detlof and Ward Edwards (1986). *Decision Analysis and Behavioral Research*. Cambridge University Press: Cambridge.
- Yi, Byeong-Uk (1992). "Rationality and the Prisoner's Dilemma in David Gauthier's *Morals by Agreement*". *The Journal of Philosophy* 89(9), pp. 484-495.

Appendices

Appendix A. CLIMEX

CLIMEX is a browser-based interactive decision-making platform that allows setting up and running a different number of economic experiments. Here we will briefly present some of the functionalities of CLIMEX, to give the reader an idea about the experiments that could be conducted using this platform. We only document here the single-player version of this platform. However, a multi-player version is now available that uses the same principles, and only adds some new features allowing the experimenter to design setups where more people contribute to the same climate fund. Before we begin, we would like to refer you back to Section 6.1.3 where the conceptual bases of the CLIMEX games is explained.

Administrator Panel

The administrator interface in CLIMEX works with four main components: games, layouts, treatments and experiments.

a) A *game* corresponds to a series of investment periods (“rounds” in Figure 1). Every period the users have to invest an amount of money from their income in a fund (see next section). The outcome of every game depends on the decisions made in the previous periods, but the games are independent from each other, i.e. in *game 2* it does not matter what you did in *game 1*.

In the “Games” panel (Figure 1), the administrator can define game templates. The shape and parameters of the function are specified here. Other settings include: the total income that players will receive every round, and whether the savings account already has some money in it or is empty.

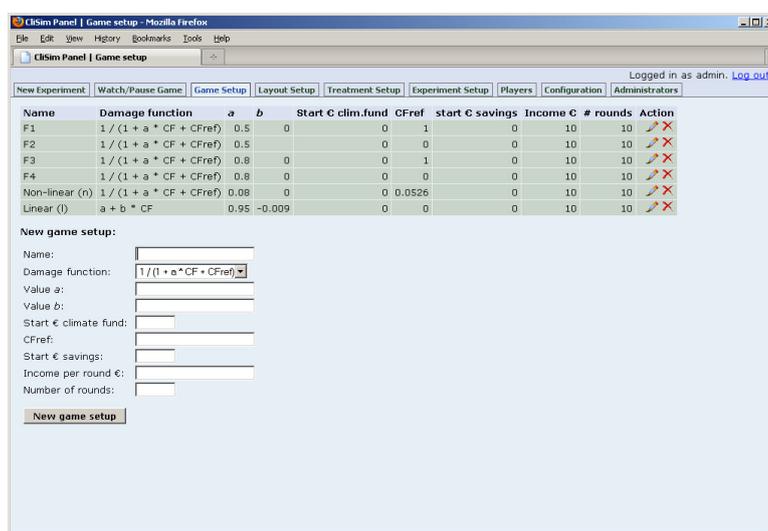


Figure 1.

b) A *layout* specifies the feedback information that the players receive about their investments during a game. The experimenter can choose whether to show or hide different elements on the screen and also define names for each of them (Figure 2). Also, the instruction screen that players will see before each game is configured here, as well as the “rules” of the game (Risk Mode).

The “Fixed Risk” option means that the output of the damage function will represent *a fraction* of the savings to be lost. Under this setting, the administrator can indicate a level of risk, which represents the probability with which the player will lose that fraction of savings. On the other hand, the “Probability of Loss” option means that the output of the damage function is a *probability* to lose **all** the savings.

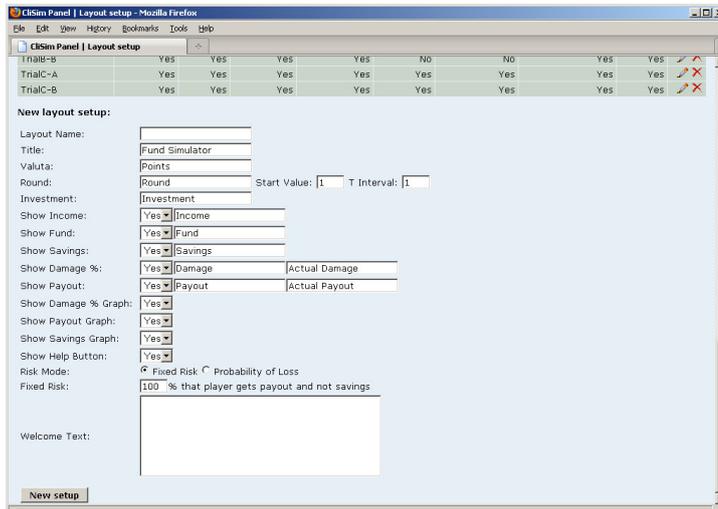


Figure 2.

c) A *treatment* can group more games that share similar characteristics. Among the settings that can be made here, we mention (see Figure 3): the type and value of noise in the information presented to the participants, the amount of delay, whether the games in that treatment are trial games or not (i.e. whether they should count for the total earnings of each player). It is also possible to specify here whether the experiment should pause after the current treatment or not. If a pause is selected, then the participants will not be able to continue to the next part of the experiment until the experimenter “un-pauses” their session from the panel. It is also in the Treatment panel that different Games are matched with a certain Layout.

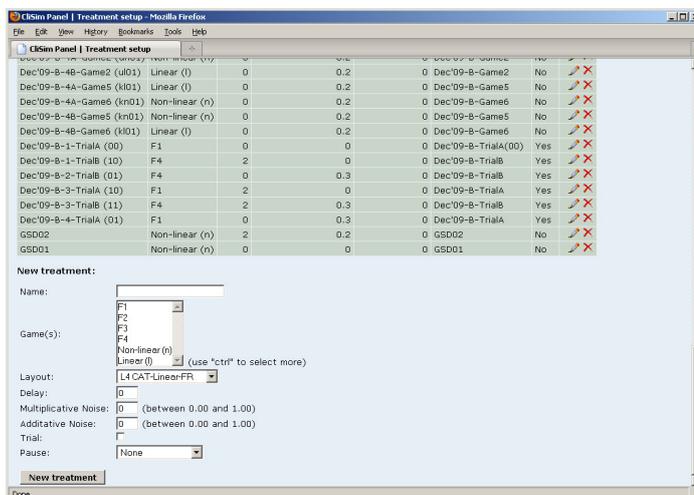


Figure 3.

d) An *experiment* is defined as a series of treatments to be played by the same participants (see Figure 4).

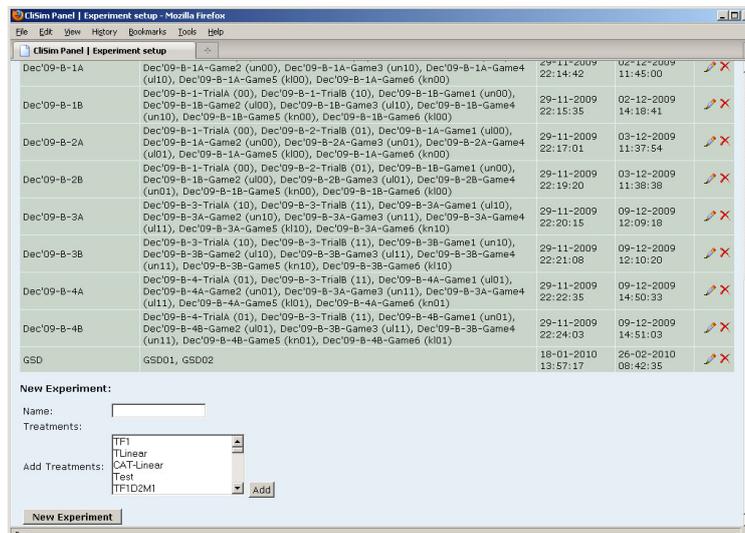


Figure 4.

The administrative panel also includes a screen where the experimenter can watch the progress of the experiment and the status of the game for each player (Figure 5). It is possible to download all the information in the database about the decisions that the players made.



Figure 5.

Finally, the “Players” screen shows the total earnings of each of the player, so that it is easy for the experimenter to prepare the payments once the experiment is over (Figure 6).

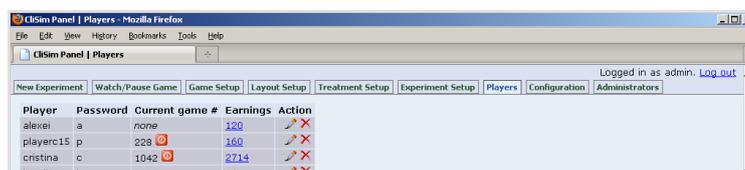


Figure 6.

User View

Each user can login by using a unique username and password that had been created for him in the administration panel (Figure 7).

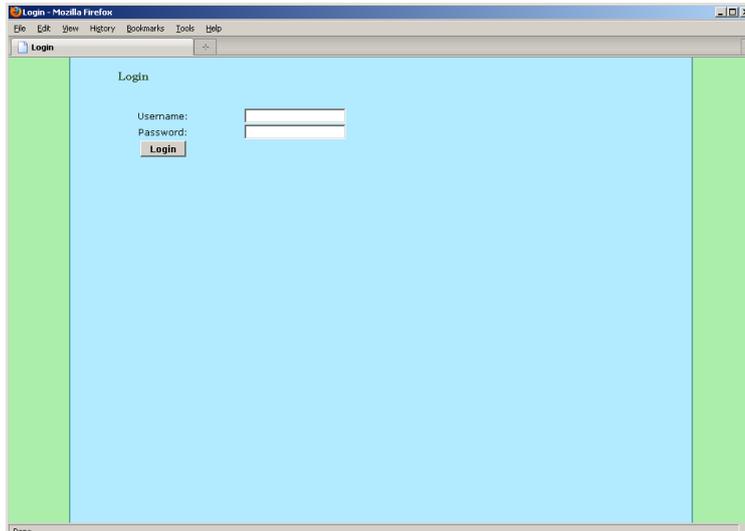


Figure 7.

First, an instruction screen is displayed (Figure 8). This information can be later on retrieved by pressing the “Help” button at any time during the experiment.

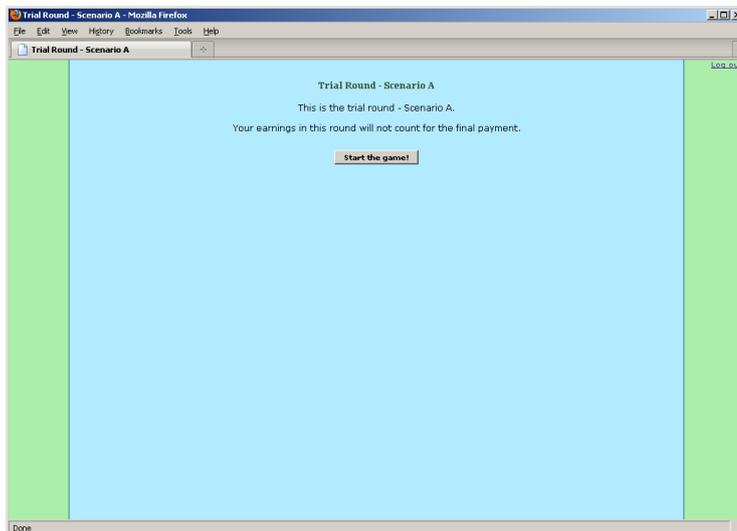


Figure 8.

Depending on the layout settings (see previous Section), the player will see one or two graphs and several rows of the decision and feedback table. He then must introduce his decision in the red square and click the “Confirm decision” button (Figure 9).

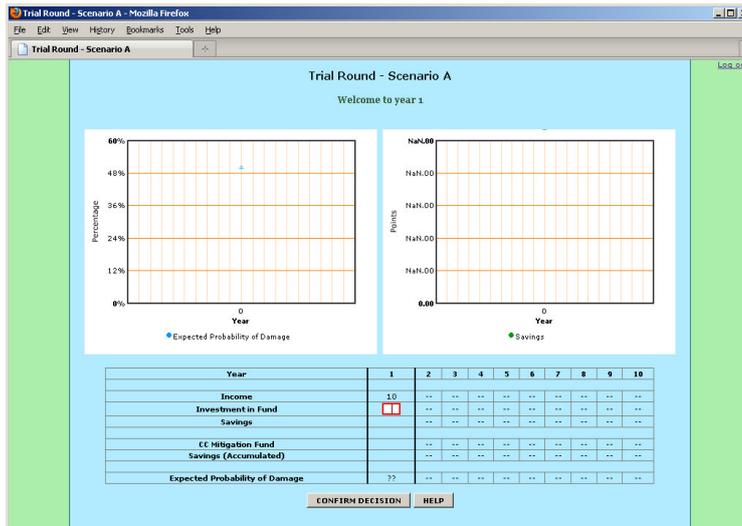


Figure 9.

Once the decision is entered, the output of the damage function is displayed in the table and in the graphs (Figure 10).

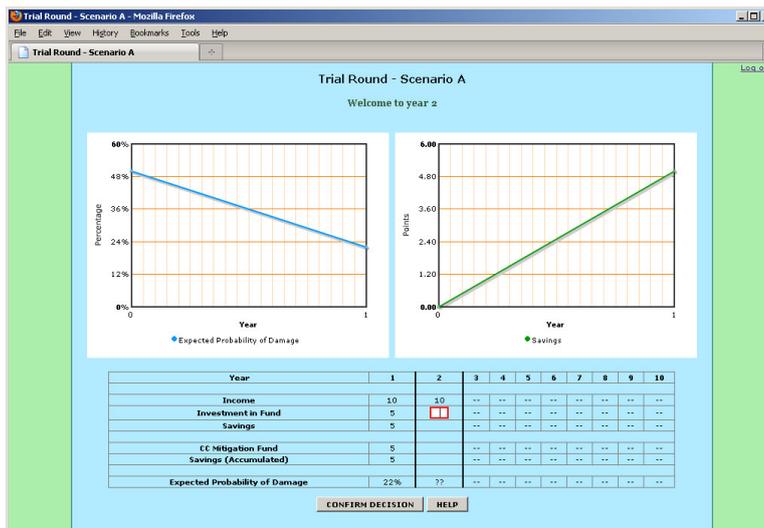


Figure 10.

At the end of each game an information screen is displayed, where a die is thrown (if that corresponds to the setup) and the total earnings of that game are explained to the player (Figure 11).

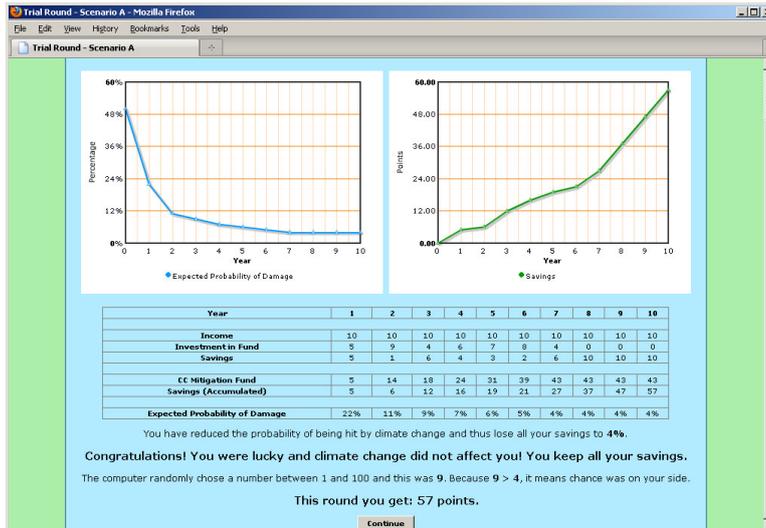


Figure 11.

Appendix B. Experiment Instructions

[as received by participants – combined here for all 4 groups]

Introduction

Welcome & Thank You

Thank you for choosing to participate in our experiment!

Please read the following instructions carefully. These instructions are the same for all participants and they explain everything you need to know for this experiment. If you have any questions, please raise your hand. One of the experimenters will approach you in order to answer your question.

Gains & Conversion

Regardless of your performance in this experiment, you will be paid a fixed amount of **2,50 Euro**.

Additionally, you can earn extra money by earning points during the experiment. The number of points that you earn depends *on you* and *on chance*. The performance and decisions of the other participants do not influence your results in any way.

The total number of points that you earn during the experiment will be exchanged at a rate of:

25 points = 1 Euro

Your performance during the experiment and your gains are confidential. The money that you earn will be paid out *in cash* at the end of the experiment without other participants being able to see how much you earned.

In order to ensure that our experiment is successful, please abide to the following simple rules:

1. Turn off your mobile phone and put it in your bag.
2. Do not communicate with other participants during the experiment.
3. Only use the functions on the screen that are necessary to carry out the experiment.

Thank you!

Overview of the Experiment

Structure

The experiment consists of 8 rounds: *two trial rounds* (not relevant for your earnings) and *six paid rounds*. At the end of these rounds, you will be asked to fill in a questionnaire.

Each round is related to one of two available *scenarios* (*scenario A and scenario B*). Further information about the scenarios will follow later. In this experiment the sequence of rounds will be as following:

Experiment Structure

Stage 0 (not relevant for your earnings):

- **one** trial round scenario A
- **one** trial round scenario B

Stage 1 (paid):

- **two** rounds scenario A
- **two** rounds scenario B

< Additional Instructions Break >

Stage 2 (paid):

- **two** rounds scenario A

<Questionnaire>

As you can see, there will be a small break between Stages 1 and 2 of the experiment, during which you will receive additional instructions for the last 2 rounds. At the end of Stage 1, the computer will automatically pause the experiment and you are asked to contact the experimenter before you can continue. Raise the coloured card (the side with “Stage 2”) that is on your desk and the experimenter will come to you. Please remain seated during the break. There won’t be any trial rounds after the break.

Keep in mind that for *every round (including the trial rounds) the experiment starts with new settings or conditions.*

Description

We will begin by briefly describing what stays the same in all rounds. Full details about how to earn points in this experiment will be presented in the section entitled “Your earnings”. Also, to help you understand the game better, some examples will follow later.

For every round there are *10 decision periods* and your earnings at the end of each round are the ones indicated after you enter your 10th decision. In other words, your final earnings depend on the situation at the end of the round.

The total number of points that you can earn is at risk because of climate change. The decisions you have to make concern and influence this risk. One decision period corresponds to one calendar year.

Every decision period (every year) you get an income of 10 points and your task is to **distribute** this income between a *personal savings account* and a *climate change mitigation fund*. You can do this by entering any value between 0 and 10 in the red box that appears in the progress table. What you contribute to the fund is marked as “Investment in Fund” and your corresponding savings that year will be calculated automatically (as Income minus the Investment in Fund) and shown in the “Savings” box (see Figure 1 below).

You should think of your income as being money for discretionary spending (i.e. it does not include the money that you need to cover your basic needs, such as food, shelter etc.).

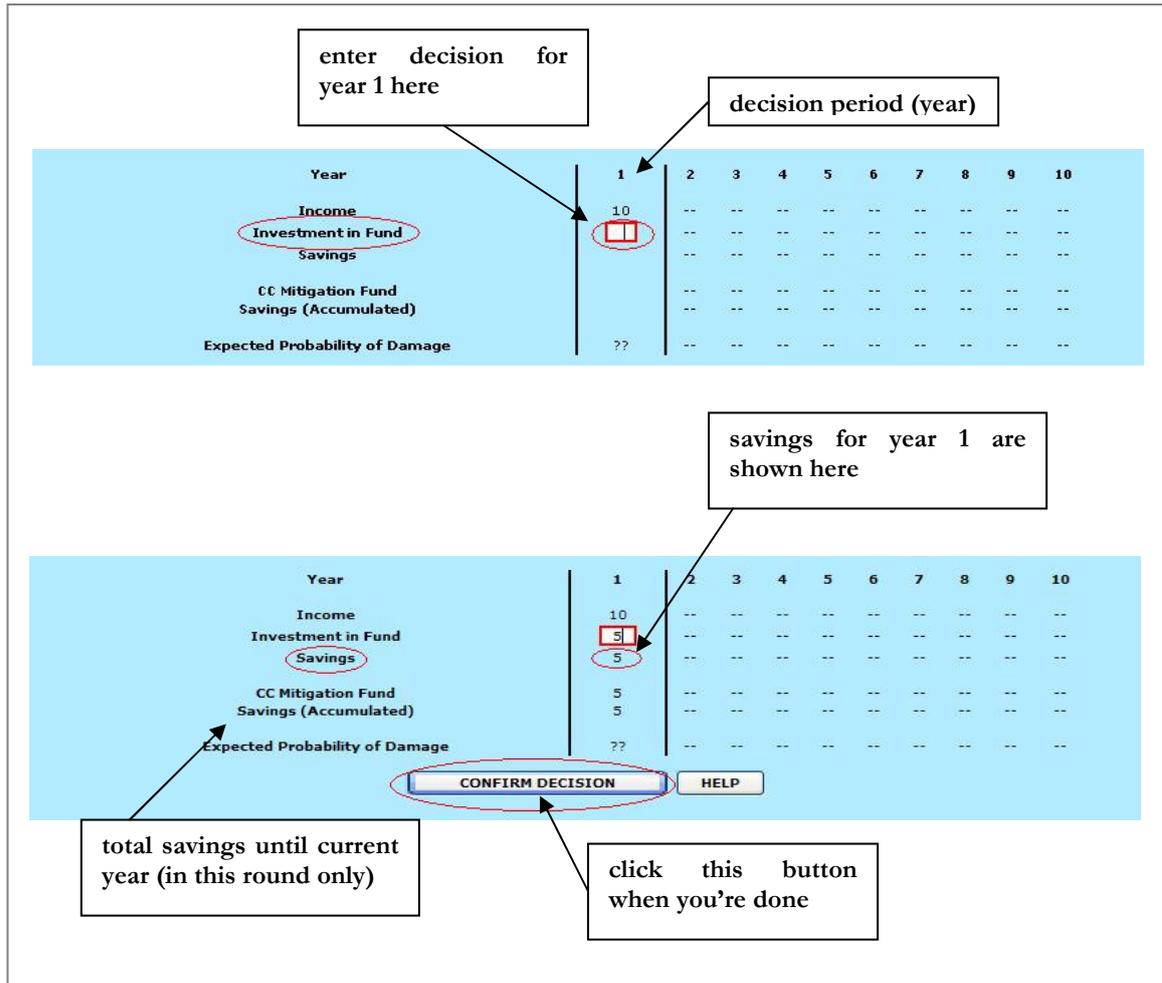


Figure 1. Elements on your screen

The points you earn in a round depend on two things:

- 1) how much you have in your *savings account* after 10 years;
- 2) whether you are *affected* by climate change or not.

More precisely, because of climate change, there is a risk that all your savings will be lost. We call this risk “*actual probability of damage*”, and we express it as a percentage, indicating the *probability that you will lose all your savings*. By investing in the *climate change mitigation fund (CC Mitigation Fund)*, you reduce this probability.

This probability is introduced to express that you do not know in which part of the world you live and whether climate change will affect you or not. If you are unlucky to be the victim of a natural disaster, for instance, you lose all your savings.

The relationship between the number of accumulated points in the climate change mitigation fund and the value of the *actual probability of damage* to your savings might change from one round to the next (including the trial rounds). In Stage 2 of the experiment you will receive additional information about this relationship.

Every round after you enter your decision and click the “Confirm decision” button, you will receive information about the newly reported *expected probability of damage* caused by climate change. Notice that this is not the same as the *actual probability of damage*. The difference will be explained later.

Your Earnings

Now we explain in more detail how the number of points that you earn depends on the decisions you make and on chance. Read the examples provided below carefully. Do not worry if you find it difficult to grasp immediately. There will be a trial round *for each scenario* to gain experience with this.

At every point in time t , the climate change **actual probability of damage** (p) is calculated as a function of the current accumulated points in the climate change mitigation fund (F):

$$p(t) = f(F(t))$$

In the first six rounds (Stage 0 and Stage 1) you do not know the function f . You will only get more information about it in Stage 2. The attributes of the function f differ from one round to another.

If the round were to end after any decision period t , the (actual) probability that you lose all your savings would be $p(t)\%$. For example, if from the formula above the software calculates that the value of p at the end of year t is 20, this means that the probability of damage is 20%. This number represents the actual chance of you being affected by climate change. If you are lucky, you will get your full savings. If not, you get nothing.

How the actual probability of damage works (example):

We start with an initial probability of damage of 60%.

In *year 1* you get an income of 10. You invest 3 points in the climate change fund and 7 in your savings. After you confirm the decision, the new value of p ($p(1)$) is calculated to be 57. This means that the *actual probability of damage* is now 57%.

In *year 2*, you receive a new income of 10. You invest 5 points in the climate change fund and 5 in your savings. This means that there are now 8 points in the fund (3+5) and 12 (7+5) points in your savings account. At the end of year 2, the new probability of damage ($p(2)$) is of 40%.

Let's assume that the game ends after year 2. You have 40% chance to lose all your savings. The computer picks a random number between 0 and 100. Let's assume this is 42. 42 is greater than 40, the value of p . This means you are lucky and climate change did not affect your savings. Your earnings this round are your full savings, i.e. 12 points.

As you can see from the example above, at the end of year 10, the computer picks a random number between 0 and 100 to see whether you were lucky or not. If it is above the value of p at the end of year 10, you earn the full savings you have accumulated until then. If it is below the value of p , you lose everything.

However, you don't know this probability exactly, as it is obscured by informational delays and distortions. The number that you see on the screen after every decision period is the *expected probability of damage*. It is not **necessarily the value of p** , but a number, that could be close to the actual probability of damage.

How the *expected* probability of damage works:

The expected probability of damage is a number that you see on the screen after every decision period which indicates what *we think* the chance of you being affected by climate change might be. This is not necessarily what this chance *really is*.

The expected probability of damage depends on the Scenario that you are in (see following section).

Please note: the values used in the example above serve only as illustration; the underlying function used here is not the function that will be used in the experiment.

Scenarios

As explained previously, the 8 rounds (including the trial ones) differ from one another in terms of the characteristics of the function f that defines the relationship between the number of points in the climate change fund and the climate change damage.

Another difference that occurs is that some rounds are in Scenario A, others in Scenario B. The basic rules are the same, what differs is the reliability of the information that you get after every round, i.e. the *expected probability of damage* depends on different things. We will now explain this difference in more detail.

[for group 1: A=No noise, no Delay, B= Delay]

Scenario A

Scenario A follows the lines of the example above. After you enter every decision you get some information about the current *probability of damage*. This information is accurate and your earnings and payment will be based on it.

Scenario B

In scenario B, we introduce delay in the information you get about the *probability of damage*.

The *delay* is set to 2 rounds and it means that although you might invest something in the climate change fund in round n , you will only see that the *probability of damage* was reduced in round $n+2$. In other words, you would not notice any improvement in the risk value until two rounds later. For instance, in the example above, after you invest 3 in round 1, you would only see the corresponding reduction in *probability of damage* from 60% to 57% in round 3.

The delay is supposed to represent uncertainties from real life. You cannot trust 100% the information you get after every decision, because some of the actions you take now may only have a visible effect later. The values that you will see in the table after every decision represent the *expected probability of damage*. However, keep in mind that the delay is just *informational* uncertainty. In the end, your earnings will depend on the *actual probability of damage* (calculated as a function of the points in the fund). Although you do not know exactly how and by which values, the *actual* value of p is immediately influenced by your actions, exactly like in Scenario A. At the end

of year 10 you will have a chance to see what the actual probabilities of damage have been during that round.

■ [end group 1]

[for group 2: A=No noise, no Delay, B= Noise]

Scenario A

Scenario A follows the lines of the example above. After you enter every decision you get some information about the current *probability of damage*. This information is accurate and your earnings and payment will be based on it.

Scenario B

In scenario B, we introduce noise in the information you get about the *probability of damage*.

The *noise* reflects an uncertainty about the value of the actual *probability of damage*. The *probability of damage* that is reported at any point t in time is not the actual one, but possibly somewhere close to it. Think of this noise aspect as representing the information distortions that are introduced by the media, for instance, or simply our insufficient understanding of the climate change phenomenon. Let's assume that the "real" *probability of damage* is indeed 57% after round 1. However, we do not know this for sure and make guesses. One year it might be reported that the climate change *probability of damage* is 58%, another year that it is 60%, then another year that it is 55%. While none of these values is the "real" one, such estimates give a good indication of the range of the "real" *probability of damage*.

The noise is supposed to represent exactly this kind of uncertainty from real life. You cannot trust 100% the information you get after every decision. The values that you see in the table after every decision represent the *estimated probability of damage*. However, keep in mind that this is just *informational* uncertainty. In the end, your earnings will depend on the *actual probability of damage* (calculated as a function of the points in the fund). Although you do not know exactly how and by which values, the *actual* value of p is immediately influenced by your actions, exactly like in Scenario A. At the end of year 10 you will have a chance to see what the *actual* probabilities of damage have been during that round.

■ [end group 2]

[for group 3: A=Delay, B= Delay and Noise]

Scenario A

In Scenario A, we introduce **delay** in the information you get about the (*actual*) *probability of damage*.

The *delay* is set to 2 periods and it means that you will only see the result of the investments in the fund that you make in year n , when you get to year $n+2$. The *expected probability of damage* that you see on the screen shows you the correct value of the *actual probability of damage*, but this happens two years later.

For instance, in the example above, after you invest 3 in period 1 (which immediately gives a reduction of the actual probability of damage to 57%), the reported *expected probability of damage* that you see on the screen will still be 60% in round 1, 60% in round 2, and finally 57% in round 3. The *expected* probability of damage in round 4

would be equal to the *actual* probability of damage from round 2 (depending on how much you invested until round 2) and so on.

The delay accounts for the real-life delays in climate and communication processes. Some of the actions you take now to mitigate climate change may only have a visible effect later. That's why we call the values that you see in the table after every decision *expected probability of damage*. This delay represents *informational* uncertainty at the moment of your decision, but in the end your earnings will depend on the *actual probability of damage* which is calculated as a function of the points in the fund. So the *actual* value of p is immediately influenced by your actions – it's just that you know it only 2 periods after your decision. At the end of year 10 the actual probabilities of damage after every decision round will be shown to you.

Scenario B

In scenario B, we not only introduce delay, but *also* **noise** in the information you get about the *probability of damage*.

The *noise* means that, after every decision period, the value of the *probability of damage* shown to you on the screen and in the table is a value slightly higher or lower than the actual value. This accounts for all kinds of distortions in our knowledge about the *actual probability of damage*. You do not know how far away it is from the “real” value.

For instance, let's assume that the *actual probability of damage* is 57% after period 1, as in the example above. If there was no delay, but there was only noise, you could see right away in the table that the probability of damage is 60%, or 52%, or any other random number that is close to 57.

Think of this *noise* aspect as representing information distortions introduced by the media, for instance, or simply our insufficient understanding of the climate change phenomenon. It is *informational* uncertainty. Although now you are even more uncertain of what the actual probability of damage might be, the *actual* value of p is immediately influenced by your actions, exactly like in Scenario A. At the end of year 10 you will have a chance to see what the *actual* probabilities of damage have been during that round.

Once more, in this scenario, the *expected probability of damage* that you see on the screen is just an estimation of the *actual* probability, as it incorporates both a delay of 2 rounds and noise.

■ [end group 3]

[for group 3: A=Noise, B= Delay and Noise]

Scenario A

In Scenario A, we introduce **noise** in the information you get about the (*actual*) *probability of damage*.

The *noise* means that, after every decision period, the value of the *probability of damage* shown to you on the screen and in the table is a value slightly higher or lower than the actual value. This accounts for all kinds of distortions in our knowledge about the *actual probability of damage*. You do not know how far away it is from the “real” value.

For instance, let's assume that the *actual probability of damage* is 57% after period 1, as in the example above. If there was no delay, but there was only noise, you could see right away in the table that the probability of damage is 60%, or 52%, or any other random number that is close to 57.

Think of this *noise* aspect as representing information distortions introduced by the media, for instance, or simply our insufficient understanding of the climate change phenomenon. It is *informational* uncertainty. Although you are uncertain of what the actual probability of damage might be, the *actual* value of p is immediately influenced by your actions. At the end of year 10 you will have a chance to see what the *actual* probabilities of damage have been during that round.

Scenario B

In scenario B, we not only introduce noise, but *also* **delay** in the information you get about the *probability of damage*.

The *delay* is set to 2 periods and it means that you will only see the result of the investments in the fund that you make in year n , when you get to year $n+2$. The *expected probability of damage* that you see on the screen shows you the correct value of the *actual probability of damage*, but this happens two years later.

For instance, let's assume that in the example above there was *no* noise, but just delay. Then, after you invest 3 in period 1 (which immediately gives a reduction of the actual probability of damage to 57%), the reported *expected probability of damage* that you see on the screen would be 60% in round 1, 60% in round 2, and finally 57% in round 3. The *expected* probability of damage in round 4 would be equal to the *actual* probability of damage from round 2 (depending on how much you invested until round 2) and so on. However, because in this scenario there is also noise, the *expected probability of damage* that you would see in this situation could look like this: 62% in round 1 (random number close to 60%), 59% in round 2 (random number close to 60%) and maybe 55% in round 3 (random number close to 57%). You will get to understand this scenario better during the trial round.

The delay accounts for the real-life delays in climate and communication processes. Some of the actions you take now to mitigate climate change may only have a visible effect later. The delay adds extra *informational* uncertainty at the moment of your decision, but in the end your earnings will depend on the *actual probability of damage* which is calculated as a function of the points in the fund. The *actual* value of p is immediately influenced by your actions, exactly like in Scenario A – it's just that you know an approximate value of it (noise) and only 2 periods after your decision (delay). At the end of year 10 the *actual* probabilities of damage after every decision round will be shown to you.

Once more, in this scenario, the *expected probability of damage* that you see on the screen is just an estimation of the *actual* probability, as it incorporates both noise and a delay of 2 rounds.

■ [end group 4]

Final Notes

Logging In

When everyone finishes reading the instructions the experimenter will announce the password that you can use to log in and start the experiment. Your username is a code of 4 characters and is indicated in the top right corner of the first page of this set of instructions.

Timing

Once we start, you may go through the rounds of the experiment at your own pace. However, you should not spend more than *10 minutes* on any round. The experimenter will tell you if you are running out of time.

Appendix C. Risk Modes

Depending on the risk discourses, the settings and rules of the game differ and so does the formalization of the payout. From the combination of collective and individual risk discourses, four possible types of game result (Figure 1):

		Collective Risk Discourse	
		Adaptation	Mitigation/Avoidance
Individual Risk Discourse	“One for all and all for one”	A	B
	“To each his own”	C	D

Figure 1. Game Risk Modes

A) The output of the damage function D represents a *fraction* of personal assets to be lost and is directly dependent on the damage.

Under this game risk mode, a participant's payout $P(T)$ after round T is:

$$P(T) = S(T) - S(T) \times D(T) = (1 - D(T)) \times S(T)$$

Everybody is affected based on the same formula, so no probabilistic differentiation is made between players.

B) The output of the damage function D represents a *probability* of loss of personal assets and affects everyone equally.

The payout $P(T)$ of each player after round T is:

$$P(T) = \begin{cases} S(T), & \text{with a probability } 1 - D(T) \\ 0, & \text{with a probability } D(T) \end{cases}$$

When played in the multiple-user version, one die is thrown once for everyone, i.e. everyone gets their full individual savings or everyone loses everything.

C) The output of the damage function D represents the *value* of the collective damage, but individual loss depends on D with a fixed probability r , $0 \leq r \leq 1$. The probability r is fixed and the same for everyone, but the actual individual loss differs:

For each player, thus, a die is separately thrown and his payout is:

$$P(T) = \begin{cases} (1 - D(T)) \times S(T), & \text{with a probability } r \\ S(T), & \text{with a probability } 1 - r \end{cases}$$

Note that on the first line the payout is calculated in the same way as in risk mode A, while the second line indicates that there is a probability that the subject will not be affected by climate change at all, and thus gain his full savings. This is probably one of the most realistic risk modes.

As a variation in the multi-user game, one die could be thrown for all players, so that everyone either gets the full savings or a fraction of it, but we leave this out as it is not an interesting option for real life.

D) The output of the damage function D is again a *probability* of loss at the end of the game, but this time it affects every player differently.

This risk mode only makes sense in the multi-player version of the game and is a variation of mode B. While $D(T)$ is the same for everyone, in this game version different dice are thrown for different individuals, i.e. some players might lose everything, while others might still get their savings.

Appendix D. Additional Instructions Sheet

[as received by participants]

You now get more information about the function f that determines the relationship between the number of points accumulated in the climate change fund (F) and the value of p at any point t in time.

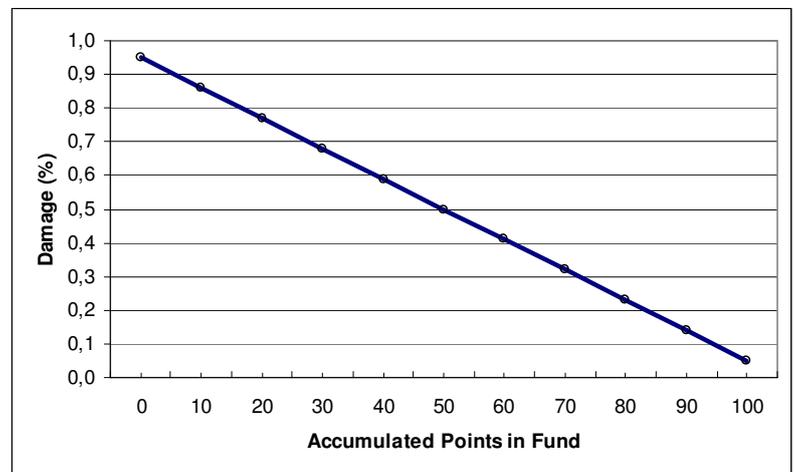
In this Stage of the experiment (Stage 2), there will be two rounds in **scenario A**:

- one with Function 1;
- one with Function 2.

The shapes and some values of the function are given below.

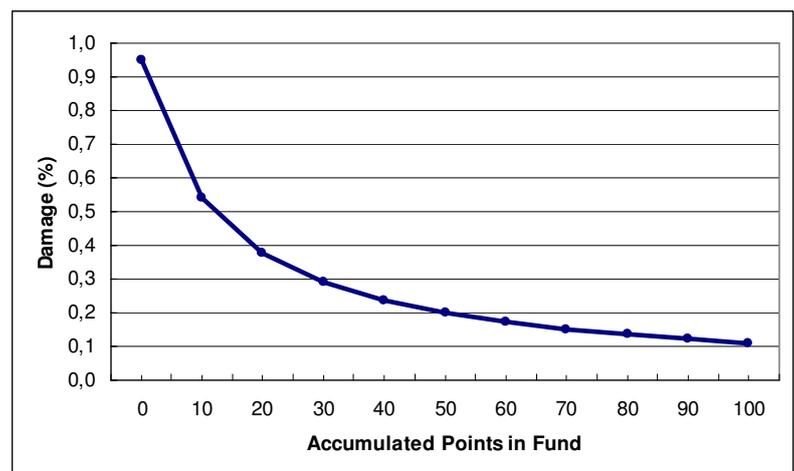
Value Table and Graph for Function 1

Accumulated points in the climate fund	Corresponding probability of damage
0	0,95
10	0,86
20	0,77
30	0,68
40	0,59
50	0,50
60	0,41
70	0,32
80	0,23
90	0,14
100	0,05



Value Table and Graph for Function 2

Accumulated points in the climate fund	Corresponding probability of damage
0	0,95
10	0,54
20	0,38
30	0,29
40	0,24
50	0,20
60	0,17
70	0,15
80	0,13
90	0,12
100	0,11



Appendix E. Questionnaire

[as received by participants]

Thank you for participating in our experiment!

We now kindly ask you to fill in this questionnaire. Please take a moment and consider all the questions carefully before giving your answer.

All your answers will stay **anonymous** and they will be solely used for the purposes of this research.

When you finish, raise the coloured card and the researcher will come to you.

Thank you!

Experiment Related Questions

The following questions are meant to help us understand what your decision strategies were during the different rounds. Please answer as accurately as possible what you were thinking *while* playing.

1. Did knowledge of the damage function play a role in your decisions?
 - Yes
 - No
2. Did you make certain decisions only to find out what would happen (“learning by experimenting”)?
 - Yes
 - No
3. Did you set an acceptable probability of damage level and then chose how much you would like to invest in the fund in total?
 - Yes
 - No
4. Did you play “casino”-style, gambling with your money?
 - Yes
 - No
5. Did you feel you could control the situation in general?
 - Yes
 - No

6. Please write in the following space any additional comments that you would like to make regarding your strategies. If you need extra space, you may use the last sheet of this questionnaire set.

Other Questions

Part A.

The following questions present you several hypothetical choices concerning amounts of money. We would like to know what choices you would make if you were in any of the situations described.

These questions will not affect in any way your earnings from the previous experiment. Also, there is no right or wrong answer to these questions.

1. Which of the following two options would you choose?
 - a. You draw a lottery ticket with a 25% chance of winning 60 euro.
 - b. You draw a lottery ticket with a 20% chance of winning 80 euro.
2. Which of the following two options would you choose?
 - a. You draw a lottery ticket with an 80% chance of winning 20 euro.
 - b. You receive 10 euro for sure.
3. We toss a coin once. Which of the following two options would you choose?
 - a. You receive 20 euro with either heads or tails;
 - b. You receive 40 euro with heads; but nothing with tails.
4. Imagine that you just won 10 euro and that you could choose between being paid in cash or receiving a lottery ticket with a certain chance to win a prize of 1000 euro. How high would that chance to win 1000 euro need to be so that you would prefer the lottery ticket to keeping the 10 euro?
5. Imagine that you just won 100 euro and that you could choose between being paid in cash or receiving a lottery ticket with a certain chance to win a prize of 1000 euro. How high would that chance to win 1000 euro need to be so that you would prefer the lottery ticket to keeping the 100 euro?
6. Imagine that you just won 500 euro and that you could choose between being paid in cash or receiving a lottery ticket with a certain chance to win a prize of 1000 euro. How high would that chance to win 1000 euro need to be so that you would prefer the lottery ticket to keeping the 500 euro?

Part B.

1. How likely do you think it is that the average annual global temperature will increase by 2°C within the following 50 years?

Very unlikely	Unlikely	Neither likely nor unlikely	Likely	Very likely
1	2	3	4	5

2. Suppose that the average annual temperature **does** increase by 2°C within the following 50 years.

For the following statements, please indicate on a 1 (very unlikely) to 5 (very likely) scale how likely you think each of the events described would be:

	Very unlikely	Unlikely	Neither likely nor unlikely	Likely	Very likely
Many people's standard of living will decrease.	1	2	3	4	5
My standard of living will decrease.	1	2	3	4	5
Starvation and food shortages will occur in many parts of the world.	1	2	3	4	5
Starvation and food shortages will occur in the area where I live.	1	2	3	4	5
Rates of diseases will increase.	1	2	3	4	5
My chances of suffering from a serious disease will increase.	1	2	3	4	5
Floods will occur in many parts of the world.	1	2	3	4	5
Parts of the Netherlands will be flooded.	1	2	3	4	5

3. Suppose the following events **do** happen. Please indicate on a 1 to 5 scale how severe do you think the negative consequences of each of these events would be:

	Almost no effects	Insignificant events	Neutral	Severe effects	Extremely severe effects
Many people's standard of living will decrease.	1	2	3	4	5
My standard of living will decrease.	1	2	3	4	5
Starvation and food shortages will occur in many parts of the world.	1	2	3	4	5
Starvation and food shortages will occur in the area where I live.	1	2	3	4	5
Rates of diseases will increase.	1	2	3	4	5
My chances of suffering from a serious disease will increase.	1	2	3	4	5
Floods will occur in many parts of the world.	1	2	3	4	5
Parts of the Netherlands will be flooded.	1	2	3	4	5

4. Suppose that the following events do happen. Please indicate what percentage of your income you **expect you would have to pay** in each of the following cases. Tick the appropriate box.

Percent of your income	<5%	5%-25%	>25%
Many people's standard of living will decrease.			
My standard of living will decrease.			
Starvation and food shortages will occur in many parts of the world.			
Starvation and food shortages will occur in the area where I live.			
Rates of diseases will increase.			
My chances of suffering from a serious disease will increase.			
Floods will occur in many parts of the world.			
Parts of the Netherlands will be flooded.			

5. Please indicate what percentage of your income you would be **willing to** invest in a climate change fund in order to avoid that any of the events listed below happens. Tick the appropriate box.

Percent of your income	<5%	5%-25%	>25%
Many people's standard of living will decrease.			
My standard of living will decrease.			
Starvation and food shortages will occur in many parts of the world.			
Starvation and food shortages will occur in the area where I live.			
Rates of diseases will increase.			
My chances of suffering from a serious disease will increase.			
Floods will occur in many parts of the world.			
Parts of the Netherlands will be flooded.			

Part C.

1. How well informed do you think you are about climate change?

Not at all informed	Somewhat uninformed		Somewhat informed	Very well informed
1	2	3	4	5

2. How relevant do you think the problem of climate change is for you?

Very irrelevant	Somewhat irrelevant		Somewhat relevant	Very relevant
1	2	3	4	5

3. How worried are you that you might be affected by climate change in the future?

Not at all worried	Not worried		Somewhat worried	Very worried
1	2	3	4	5

4. How would you evaluate your own risk attitude?

Please indicate whether you agree with the following statements or not:

	Strongly disagree	Somewhat disagree	Neither agree, nor disagree	Somewhat agree	Strongly agree	
The government will inform me if there is a serious risk that I could be affected by climate change.	1	2	3	4	5	
When reporting on climate change issues mass-media is usually neutral.	1	2	3	4	5	

Part D.

Please fill in the following fields about you:

Gender: female ; male .

Age: _____

Occupation: _____

Higher degree obtained so far: _____

Study/Working field: _____

Are you a member of an environmental organization?

- yes
- no

Income level:

- less than 1500
- between 1500 and 2000
- 2500-3500
- 3500-4500
- >4500

Appendix F. Optimum Investments

Linear Function – Optimum Investment & Maximum Expected Payout

$$D(F) = a + bF$$

$$a = 0.95$$

$$b = -0.009$$

$$I = 100$$

$$EP = (I - F) \cdot (1 - D(F))$$

$$EP \text{ max} \Rightarrow EP' = 0$$

$$EP = (I - F)(1 - a - bF) = I - aI - bFI - F + aF + bF^2$$

$$EP' = 2bF + a - bI - 1$$

$$EP' = 0 \Rightarrow 2bF + a - bI - 1 = 0$$

$$-2 \cdot 0.009F + 0.95 + 0.009 \cdot 100 - 1 = 0$$

$$0.018F = 1.90 - 0.95$$

$$F_0 = F = \frac{0.85}{0.018} \approx 47.2$$

$$D(47.2) = 0.95 - 0.009 \cdot 47.2 \approx 0.525$$

$$EP_{\text{max}} = (100 - 47.2) \cdot 0.475 = 25.08$$

Non-Linear Function – Optimum Investment & Maximum Expected Payout

$$D(F) = \frac{1}{a + bF}$$

$$a = 1.0526$$

$$b = 0.08$$

$$a, b > 0$$

$$I = 100$$

$$EP = (I - F) \cdot (1 - D(F))$$

$$EP \text{ max} \Rightarrow P' = 0$$

$$EP = (I - F) \left(1 - \frac{1}{a + bF}\right) = I - F - \frac{I}{a + bF} + \frac{F}{a + bF}$$

$$EP' = -1 - I \left(\frac{1}{a + bF}\right)' + \left(\frac{F}{a + bF}\right)' = -1 + I \left(\frac{(a + bF)'}{(a + bF)^2}\right) + \left(\frac{F}{a + bF}\right)' =$$

$$= -1 + \frac{Ib}{(a + bF)^2} + \frac{F'}{a + bF} - \frac{F(a + bF)'}{(a + bF)^2} = -1 + \frac{Ib}{(a + bF)^2} + \frac{1}{a + bF} - \frac{bF}{(a + bF)^2}$$

$$EP' = 0$$

$$\Rightarrow -1 + \frac{bI - bF}{(a + bF)^2} + \frac{1}{a + bF} = 0$$

$$-(a + bF)^2 + bI - bF + a + bF = 0$$

$$-a^2 - 2abF - b^2F^2 + bI + a = 0$$

$$-b^2F^2 - 2abF + (bI + a - a^2) = 0$$

$$\Delta = 4a^2b^2 + 4b^2(bI + a - a^2) = 4b^3I + 4ab^2 = 4b^2(bI + a)$$

$$F_{1,2} = -\frac{2ab \pm \sqrt{4b^2(bI + a)}}{2b^2} = -\frac{a \pm \sqrt{(bI + a)}}{b}$$

$$\Rightarrow F_{1,2} = -\frac{1.0526 \cdot 0.08 \pm \sqrt{0.08 \cdot 100 + 1.0526}}{0.08} = -\frac{1.0526 \pm \sqrt{9.0526}}{0.08}$$

$$F_1 = -\frac{1.0526 - 3.0087}{0.08} = 24.4519 \approx 24.5$$

$$F_2 \approx -\frac{1.0526 + 3.0087}{0.08} = -50.766$$

We need $F_0 \geq 0$, thus $F_0 = F_1$:

$$D(24.5) = \frac{1}{1.0526 + 0.08 \cdot 24.5} \approx 0.33$$

$$P_{\text{max}} = (100 - 24.5) \cdot 0.67 \approx 50.6$$

Appendix G. Random Variables

In order to calculate the mean and standard deviation of a random variable, the following formulas are used:

$$\mu = \sum_{i=1}^n p_i \cdot o_i$$

$$\text{var}(x) = \sum_{i=1}^n p_i \cdot (o_i - \mu)^2$$

$$\sigma = \sqrt{\text{var}(x)}$$

where p_i is the probability of winning the amount o_i and i represents the possible outcomes.

E.g. For the bet “60 euro with probability 25%” (60:0.25), it is implied that there is 75% of getting 0 and we make the following calculations:

$$EV = \mu = 60 \times 0.25 + 0 \times 0.75 = 15$$

$$\text{var} = 0.25 \times (60 - 15)^2 + 0.75 \times (0 - 15)^2 = 0.25 \times 2025 + 0.75 \times 225 = 506.25 + 168.75 = 675$$

$$\sigma = \sqrt{675} = 25.98$$

Appendix H. Statistical Tables

List of Procedures w/ Appendix H Page Number:

1) UNIANOVA on CFInv, LearningEffect.....	129
2) One-Way ANOVA on NormDevOpt: UNCTYTotal as Factor, File Split by Malignancy	130
3) UNIANOVA on NormDevOpt.....	133
4) UNIANOVA on NormDevOpt, File Split by Malignancy	136
5) Test of Homogeneity of Regression Assumption for Covariates, with CFInv as Dependent Variable.....	138
6) UNIANCOVA on CFInv (excluded ExpEff, Worry, Info, Likelihood and Relevance).....	139
7) Curve Fit: Risk Score 2->Final Damage Level, File Split by Malignancy	148

1) UNIANOVA on CFInv, LearningEffect

Between-Subjects Factors			
		Value Label	N
Uncertainty	000	Baseline (K)	80
	001	Noise (K)	40
	010	Delay (K)	40
	100	Baseline (U)	80
	101	Noise (U)	80
	110	Delay (U)	80
	111	Delay&Noise (U)	80
	Function shape	1	Linear
2		Non-linear	240
Game number	1		80
	2		80
	3		80
	4		80
	5		80
	6		80

Tests of Between-Subjects Effects

Dependent Variable: Total investment

Source	Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power ^b
Corrected Model	30070.981 ^a	35	859.171	4.041	.000	.242	141.435	1.000
Intercept	1123872.499	1	1123872.499	5285.996	.000	.923	5285.996	1.000
UNCTYTotal	342.669	5	68.534	.322	.900	.004	1.612	.132
Malignancy	22498.615	1	22498.615	105.819	.000	.192	105.819	1.000
Learningeff	555.718	4	138.929	.653	.625	.006	2.614	.213
UNCTYTotal * Malignancy	421.069	5	84.214	.396	.852	.004	1.980	.154
UNCTYTotal * Learningeff	116.669	7	16.667	.078	.999	.001	.549	.070

Malignancy * Learningeff	1455.847	4	363.962	1.712	.146	.015	6.847	.525
UNCTYTotal * Malignancy * Learningeff	2304.769	7	329.253	1.549	.149	.024	10.840	.648
Error	94400.250	444	212.613					
Total	1306241.000	480						
Corrected Total	124471.231	479						

a. R Squared = .242 (Adjusted R Squared = .182)

b. Computed using alpha = .05

2) One-Way ANOVA on NormDevOpt: UNCTYTotal as Factor, File Split by Malignancy

Test of Homogeneity of Variances

Normalized Deviation from Optimum

Function shape	Levene Statistic	df1	df2	Sig.
Linear	1.337	6	233	.241
Non-linear	.418	6	233	.867

ANOVA

Normalized Deviation from Optimum

Function shape		Sum of Squares	df	Mean Square	F	Sig.
Linear	Between Groups	.844	6	.141	1.463	.192
	Within Groups	22.413	233	.096		
	Total	23.257	239			
Non-linear	Between Groups	.940	6	.157	.446	.848
	Within Groups	81.918	233	.352		
	Total	82.858	239			

Multiple Comparisons

Normalized Deviation from Optimum

Bonferroni

Function shape	(I) UNCTYNum	(J) UNCTYNum	Mean Difference (I-J)	Std. Error	Sig.	95% Confidence Interval	
						Lower Bound	Upper Bound
Linear	Baseline (K)	Noise (K)	-.03602	.08494	1.000	-.2969	.2249
		Delay (K)	-.06568	.08494	1.000	-.3266	.1952
		Baseline (U)	-.02489	.06935	1.000	-.2379	.1881
		Noise (U)	.08263	.06935	1.000	-.1304	.2957
		Delay (U)	.04449	.06935	1.000	-.1685	.2575

		Delay & Noise (U)	.11653	.06935	1.000	-.0965	.3296
Noise (K)		Baseline (K)	.03602	.08494	1.000	-.2249	.2969
		Delay (K)	-.02966	.09808	1.000	-.3309	.2716
		Baseline (U)	.01112	.08494	1.000	-.2498	.2720
		Noise (U)	.11864	.08494	1.000	-.1423	.3796
		Delay (U)	.08051	.08494	1.000	-.1804	.3414
		Delay & Noise (U)	.15254	.08494	1.000	-.1084	.4135
Delay (K)		Baseline (K)	.06568	.08494	1.000	-.1952	.3266
		Noise (K)	.02966	.09808	1.000	-.2716	.3309
		Baseline (U)	.04078	.08494	1.000	-.2201	.3017
		Noise (U)	.14831	.08494	1.000	-.1126	.4092
		Delay (U)	.11017	.08494	1.000	-.1507	.3711
		Delay & Noise (U)	.18220	.08494	.693	-.0787	.4431
Baseline (U)		Baseline (K)	.02489	.06935	1.000	-.1881	.2379
		Noise (K)	-.01112	.08494	1.000	-.2720	.2498
		Delay (K)	-.04078	.08494	1.000	-.3017	.2201
		Noise (U)	.10752	.06935	1.000	-.1055	.3206
		Delay (U)	.06939	.06935	1.000	-.1436	.2824
		Delay & Noise (U)	.14142	.06935	.894	-.0716	.3545
Noise (U)		Baseline (K)	-.08263	.06935	1.000	-.2957	.1304
		Noise (K)	-.11864	.08494	1.000	-.3796	.1423
		Delay (K)	-.14831	.08494	1.000	-.4092	.1126
		Baseline (U)	-.10752	.06935	1.000	-.3206	.1055
		Delay (U)	-.03814	.06935	1.000	-.2512	.1749
		Delay & Noise (U)	.03390	.06935	1.000	-.1791	.2469
Delay (U)		Baseline (K)	-.04449	.06935	1.000	-.2575	.1685
		Noise (K)	-.08051	.08494	1.000	-.3414	.1804
		Delay (K)	-.11017	.08494	1.000	-.3711	.1507
		Baseline (U)	-.06939	.06935	1.000	-.2824	.1436
		Noise (U)	.03814	.06935	1.000	-.1749	.2512
		Delay & Noise (U)	.07203	.06935	1.000	-.1410	.2851
Delay & Noise (U)		Baseline (K)	-.11653	.06935	1.000	-.3296	.0965
		Noise (K)	-.15254	.08494	1.000	-.4135	.1084
		Delay (K)	-.18220	.08494	.693	-.4431	.0787
		Baseline (U)	-.14142	.06935	.894	-.3545	.0716
		Noise (U)	-.03390	.06935	1.000	-.2469	.1791
		Delay (U)	-.07203	.06935	1.000	-.2851	.1410
Non-linear	Baseline (K)	Noise (K)	-.08673	.16238	1.000	-.5855	.4121
		Delay (K)	-.07449	.16238	1.000	-.5733	.4243
		Baseline (U)	-.12449	.13259	1.000	-.5318	.2828
		Noise (U)	-.16429	.13259	1.000	-.5716	.2430

	Delay (U)	-.05000	.13259	1.000	-.4573	.3573
	Delay & Noise (U)	-.17653	.13259	1.000	-.5838	.2307
Noise (K)	Baseline (K)	.08673	.16238	1.000	-.4121	.5855
	Delay (K)	.01224	.18750	1.000	-.5637	.5882
	Baseline (U)	-.03776	.16238	1.000	-.5366	.4610
	Noise (U)	-.07755	.16238	1.000	-.5764	.4213
	Delay (U)	.03673	.16238	1.000	-.4621	.5355
	Delay & Noise (U)	-.08980	.16238	1.000	-.5886	.4090
Delay (K)	Baseline (K)	.07449	.16238	1.000	-.4243	.5733
	Noise (K)	-.01224	.18750	1.000	-.5882	.5637
	Baseline (U)	-.05000	.16238	1.000	-.5488	.4488
	Noise (U)	-.08980	.16238	1.000	-.5886	.4090
	Delay (U)	.02449	.16238	1.000	-.4743	.5233
	Delay & Noise (U)	-.10204	.16238	1.000	-.6008	.3968
Baseline (U)	Baseline (K)	.12449	.13259	1.000	-.2828	.5318
	Noise (K)	.03776	.16238	1.000	-.4610	.5366
	Delay (K)	.05000	.16238	1.000	-.4488	.5488
	Noise (U)	-.03980	.13259	1.000	-.4471	.3675
	Delay (U)	.07449	.13259	1.000	-.3328	.4818
	Delay & Noise (U)	-.05204	.13259	1.000	-.4593	.3552
Noise (U)	Baseline (K)	.16429	.13259	1.000	-.2430	.5716
	Noise (K)	.07755	.16238	1.000	-.4213	.5764
	Delay (K)	.08980	.16238	1.000	-.4090	.5886
	Baseline (U)	.03980	.13259	1.000	-.3675	.4471
	Delay (U)	.11429	.13259	1.000	-.2930	.5216
	Delay & Noise (U)	-.01224	.13259	1.000	-.4195	.3950
Delay (U)	Baseline (K)	.05000	.13259	1.000	-.3573	.4573
	Noise (K)	-.03673	.16238	1.000	-.5355	.4621
	Delay (K)	-.02449	.16238	1.000	-.5233	.4743
	Baseline (U)	-.07449	.13259	1.000	-.4818	.3328
	Noise (U)	-.11429	.13259	1.000	-.5216	.2930
	Delay & Noise (U)	-.12653	.13259	1.000	-.5338	.2807
Delay & Noise (U)	Baseline (K)	.17653	.13259	1.000	-.2307	.5838
	Noise (K)	.08980	.16238	1.000	-.4090	.5886
	Delay (K)	.10204	.16238	1.000	-.3968	.6008
	Baseline (U)	.05204	.13259	1.000	-.3552	.4593
	Noise (U)	.01224	.13259	1.000	-.3950	.4195
	Delay (U)	.12653	.13259	1.000	-.2807	.5338

3) UNIANOVA on NormDevOptⁱ

Between-Subjects Factors

		Value Label	N
Uncertainty	000	Baseline (K)	80
	001	Noise (K)	40
	010	Delay (K)	40
	100	Baseline (U)	80
	101	Noise (U)	80
	110	Delay (U)	80
	111	Delay&Noise (U)	80
	Function shape	1	Linear
2		Non-linear	240
Order	A		240
	B		240

Tests of Between-Subjects Effects

Dependent Variable: Dist_Opt_Norm

Source	Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power ^b
Corrected Model	42.351 ^a	27	1.569	7.127	.000	.299	192.427	1.000
Intercept	98.156	1	98.156	445.988	.000	.497	445.988	1.000
UNCTYTotal	.389	6	.065	.294	.940	.004	1.767	.132
Malignancy ⁱⁱ	30.773	1	30.773	139.823	.000	.236	139.823	1.000
Order ⁱⁱⁱ	2.034	1	2.034	9.240	.003	.020	9.240	.858
UNCTYTotal * Malignancy	1.396	6	.233	1.057	.388	.014	6.341	.420
UNCTYTotal * Order	1.632	6	.272	1.236	.287	.016	7.414	.488
Malignancy * Order	1.028	1	1.028	4.671	.031	.010	4.671	.578
UNCTYTotal * Malignancy * Order	.122	6	.020	.092	.997	.001	.554	.073

ⁱ In the statistical tables NormDevOpt sometimes appears as Dist_Opt_Norm.

ⁱⁱ This effect is of course significant, due to the intrinsic characteristics of the *NormDevOpt* measure (see Section 8.2.1 for details). We ignore it here.

ⁱⁱⁱ The order effect might not be significant here, but also seem significant due to the differences that exist in the measures of *NormDevOpt*. This would also explain the interaction Malignancy*Order. That is why it is better to test for Order effects by running the analysis of variance on the file split by malignancy (see next section).

Error	99.479	452	.220				
Total	248.054	480					
Corrected Total	141.830	479					

a. R Squared = .299 (Adjusted R Squared = .257)

b. Computed using alpha = .05

Uncertainty (not significant experiment-wide):

Pairwise Comparisons

Dependent Variable: Dist_Opt_Norm

(I) Uncertainty	(J) Uncertainty	Mean Difference (I-J)	Std. Error	Sig. ^a	95% Confidence Interval for Difference ^a	
					Lower Bound	Upper Bound
Baseline (K)	Noise (K)	-.061	.091	.500	-.240	.117
	Delay (K)	-.070	.091	.441	-.249	.108
	Baseline (U)	-.075	.074	.314	-.220	.071
	Noise (U)	-.041	.074	.582	-.187	.105
	Delay (U)	-.003	.074	.970	-.149	.143
	Delay&Noise (U)	-.030	.074	.686	-.176	.116
Noise (K)	Baseline (K)	.061	.091	.500	-.117	.240
	Delay (K)	-.009	.105	.934	-.215	.197
	Baseline (U)	-.013	.091	.884	-.192	.165
	Noise (U)	.021	.091	.821	-.158	.199
	Delay (U)	.059	.091	.519	-.120	.237
	Delay&Noise (U)	.031	.091	.730	-.147	.210
Delay (K)	Baseline (K)	.070	.091	.441	-.108	.249
	Noise (K)	.009	.105	.934	-.197	.215
	Baseline (U)	-.005	.091	.960	-.183	.174
	Noise (U)	.029	.091	.748	-.149	.208
	Delay (U)	.067	.091	.459	-.111	.246
	Delay&Noise (U)	.040	.091	.659	-.138	.219
Baseline (U)	Baseline (K)	.075	.074	.314	-.071	.220
	Noise (K)	.013	.091	.884	-.165	.192
	Delay (K)	.005	.091	.960	-.174	.183
	Noise (U)	.034	.074	.648	-.112	.180
	Delay (U)	.072	.074	.333	-.074	.218
	Delay&Noise (U)	.045	.074	.547	-.101	.190
Noise (U)	Baseline (K)	.041	.074	.582	-.105	.187
	Noise (K)	-.021	.091	.821	-.199	.158
	Delay (K)	-.029	.091	.748	-.208	.149
	Baseline (U)	-.034	.074	.648	-.180	.112
	Delay (U)	.038	.074	.608	-.108	.184
	Delay&Noise (U)	.011	.074	.884	-.135	.157
Delay (U)	Baseline (K)	.003	.074	.970	-.143	.149
	Noise (K)	-.059	.091	.519	-.237	.120
	Delay (K)	-.067	.091	.459	-.246	.111

	Baseline (U)	-.072	.074	.333	-.218	.074
	Noise (U)	-.038	.074	.608	-.184	.108
	Delay&Noise (U)	-.027	.074	.714	-.173	.119
Delay&Noise (U)	Baseline (K)	.030	.074	.686	-.116	.176
	Noise (K)	-.031	.091	.730	-.210	.147
	Delay (K)	-.040	.091	.659	-.219	.138
	Baseline (U)	-.045	.074	.547	-.190	.101
	Noise (U)	-.011	.074	.884	-.157	.135
	Delay (U)	.027	.074	.714	-.119	.173

Based on estimated marginal means

a. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

Function shape (significant, but this might be due to the opportunity for higher investments – see discussion in Section 7.2.1):

Pairwise Comparisons

Dependent Variable:Dist_Opt_Norm

(I) Function shape	(J) Function shape	Mean Difference (I-J)	Std. Error	Sig. ^a	95% Confidence Interval for Difference ^a	
					Lower Bound	Upper Bound
Linear	Non-linear	-.532 [*]	.045	.000	-.620	-.443
Non-linear	Linear	.532 [*]	.045	.000	.443	.620

Based on estimated marginal means

*. The mean difference is significant at the .05 level.

a. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

Order (significant):

Pairwise Comparisons

Dependent Variable:Dist_Opt_Norm

(I) Order	(J) Order	Mean Difference (I-J)	Std. Error	Sig. ^a	95% Confidence Interval for Difference ^a	
					Lower Bound	Upper Bound
A	B	-.137 [*]	.045	.003	-.225	-.048
B	A	.137 [*]	.045	.003	.048	.225

Based on estimated marginal means

*. The mean difference is significant at the .05 level.

a. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

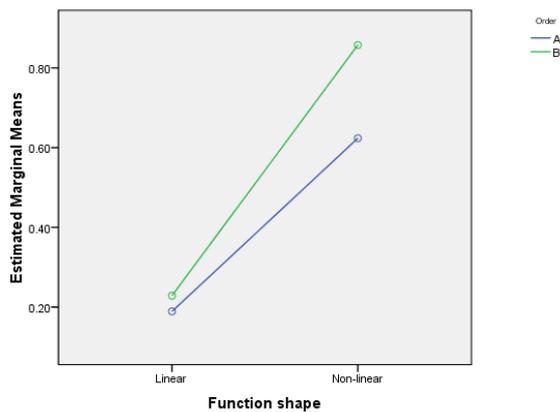
Malignancy * Order (significant):

4. Function shape * Order

Dependent Variable:Dist_Opt_Norm

Function shape	Order	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
Linear	A	.189	.045	.101	.278
	B	.229	.045	.140	.317
Non-linear	A	.624	.045	.535	.712
	B	.857	.045	.769	.946

Estimated Marginal Means of Dist_Opt_Norm



4) UNIANOVA on NormDevOpt, File Split by Malignancy

Between-Subjects Factors

Function shape			Value Label	N
Linear	Order	A		120
		B		120
	Uncertainty	000	Baseline (K)	40
		001	Noise (K)	20
		010	Delay (K)	20
		100	Baseline (U)	40
		101	Noise (U)	40
		110	Delay (U)	40
		111	Delay&Noise (U)	40
Non-linear	Order	A		120
		B		120
	Uncertainty	000	Baseline (K)	40
		001	Noise (K)	20
		010	Delay (K)	20
		100	Baseline (U)	40
		101	Noise (U)	40
		110	Delay (U)	40
		111	Delay&Noise (U)	40

Tests of Between-Subjects Effects

Dependent Variable:Normalized Deviation from Optimum

Function shape	Source	Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power ^b
Linear	Corrected Model	1.464 ^a	13	.113	1.167	.305	.063	15.177	.687
	Intercept	9.505	1	9.505	98.565	.000	.304	98.565	1.000
	Order	.085	1	.085	.881	.349	.004	.881	.154
	UNCTYTotal	.844	6	.141	1.459	.193	.037	8.757	.562
	Order * UNCTYTotal	.563	6	.094	.973	.444	.025	5.840	.382
	Error	21.794	226	.096					
	Total	32.633	240						
	Corrected Total	23.257	239						
Non-linear	Corrected Model	5.173 ^c	13	.398	1.158	.312	.062	15.048	.682
	Intercept	119.424	1	119.424	347.425	.000	.606	347.425	1.000
	Order	2.977	1	2.977	8.659	.004	.037	8.659	.834
	UNCTYTotal	.940	6	.157	.456	.840	.012	2.735	.185
	Order * UNCTYTotal	1.191	6	.198	.577	.748	.015	3.464	.229
	Error	77.685	226	.344					
	Total	215.420	240						
	Corrected Total	82.858	239						

a. R Squared = .063 (Adjusted R Squared = .009)

b. Computed using alpha = .05

c. R Squared = .062 (Adjusted R Squared = .008)

Order (significant for the Non-linear treatment):

Order

Dependent Variable:Normalized Deviation from Optimum

Function shape	Order	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
Linear	A	.189	.030	.131	.248
	B	.229	.030	.170	.287
Non-linear	A	.624	.056	.513	.734
	B	.857	.056	.747	.968

5) Test of Homogeneity of Regression Assumption for Covariates, with CFInv as Dependent Variable

Tests of Between-Subjects Effects

Dependent Variable: Dist_Opt_Norm

Source	Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power ^b
Corrected Model	61.522 ^a	72	.854	4.330	.000	.434	311.793	1.000
Intercept	1.327	1	1.327	6.728	.010	.016	6.728	.735
Malignancy * RiskScore2	.284	1	.284	1.438	.231	.004	1.438	.223
Order * RiskScore2	.346	1	.346	1.755	.186	.004	1.755	.262
UNCTYTotal * RiskScore2	1.864	6	.311	1.574	.153	.023	9.446	.606
Malignancy * Likelihood	.000	1	.000	.001	.977	.000	.001	.050
Order * Likelihood	.024	1	.024	.120	.730	.000	.120	.064
UNCTYTotal * Likelihood	.453	6	.076	.383	.890	.006	2.297	.161
Malignancy * Relevance	.049	1	.049	.248	.619	.001	.248	.079
Order * Relevance	.019	1	.019	.098	.755	.000	.098	.061
UNCTYTotal * Relevance	1.072	6	.179	.905	.491	.013	5.431	.359
Malignancy * Information	.205	1	.205	1.039	.309	.003	1.039	.174
Order * Information	.553	1	.553	2.801	.095	.007	2.801	.386
UNCTYTotal * Information	1.540	6	.257	1.301	.255	.019	7.807	.511
Malignancy * ^{iv} Worry	1.386	1	1.386	7.026	.008	.017	7.026	.753
Order * Worry	.041	1	.041	.206	.650	.001	.206	.074
UNCTYTotal * Worry	1.044	6	.174	.882	.508	.013	5.290	.350
Malignancy * ExpEff	.025	1	.025	.129	.720	.000	.129	.065
Order * ExpEff	.407	1	.407	2.061	.152	.005	2.061	.299
UNCTYTotal * ExpEff	.499	6	.083	.421	.865	.006	2.526	.175
Malignancy * GovTrust	.426	1	.426	2.157	.143	.005	2.157	.311
Order * GovTrust	.511	1	.511	2.588	.108	.006	2.588	.361
UNCTYTotal * GovTrust	.907	6	.151	.766	.597	.011	4.594	.304
Malignancy * MedTrust	.004	1	.004	.018	.892	.000	.018	.052
Order * MedTrust	.822	1	.822	4.163	.042	.010	4.163	.530

^{iv} Covariates that have a significant interaction with one of the independent variables will be eliminated from further analysis.

UNCTYTotal * MedTrust	.747	6	.125	.631	.705	.009	3.786	.252
Error	80.308	407	.197					
Total	248.054	480						
Corrected Total	141.830	479						

a. R Squared = .434 (Adjusted R Squared = .334)

b. Computed using alpha = .05

6) UNIANCOVA on CFinv (excluded ExpEff, Worry, Info, Likelihood and Relevance)^v

Between-Subjects Factors

	Value	Label	N
Uncertainty	000	Baseline (K)	80
	001	Noise (K)	40
	010	Delay (K)	40
	100	Baseline (U)	80
	101	Noise (U)	80
	110	Delay (U)	80
	111	Delay&Noise (U)	80
	Function shape	1	Linear
2		Non-linear	240
Order	A		240
	B		240

Levene's Test of Equality of Error Variances^a

Dependent Variable: Total Investment

F	df1	df2	Sig.
.721	27	452	.848

Tests the null hypothesis that the error variance of the dependent variable is equal across groups.

a. Design: Intercept + RiskScore2 + GovTrust + MedTrust + UNCTYTotal + Malignancy + Order + UNCTYTotal * Malignancy + UNCTYTotal * Order + Malignancy * Order + UNCTYTotal * Malignancy * Order

^v These covariates were excluded either because they did not satisfy the homogeneity of regression assumption, or because they turned out to have an insignificant effect when we first ran the model with them.

Tests of Between-Subjects Effects

Dependent Variable: Total Investment

Source	Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power ^b
Corrected Model	35988.410 ^a	30	1199.614	6.087	.000	.289	182.621	1.000
Intercept	63352.510	1	63352.510	321.478	.000	.417	321.478	1.000
RiskScore2	2537.726	1	2537.726	12.878	.000	.028	12.878	.947
GovTrust	2130.080	1	2130.080	10.809	.001	.024	10.809	.907
MedTrust	1840.508	1	1840.508	9.340	.002	.020	9.340	.862
UNCTYTotal	659.605	6	109.934	.558	.764	.007	3.347	.225
Malignancy	22635.235	1	22635.235	114.861	.000	.204	114.861	1.000
Order	1929.481	1	1929.481	9.791	.002	.021	9.791	.877
UNCTYTotal * Malignancy	1804.960	6	300.827	1.527	.168	.020	9.159	.591
UNCTYTotal * Order	1403.580	6	233.930	1.187	.312	.016	7.122	.470
Malignancy * Order	406.501	1	406.501	2.063	.152	.005	2.063	.300
UNCTYTotal * Malignancy * Order	111.644	6	18.607	.094	.997	.001	.567	.073
Error	88482.822	449	197.066					
Total	1306241.000	480						
Corrected Total	124471.231	479						

a. R Squared = .289 (Adjusted R Squared = .242)

b. Computed using alpha = .05

Parameter Estimates

Dependent Variable: Total Investment

Parameter	B	Std. Error	t	Sig.	95% Confidence Interval		Partial Eta Squared	Noncent. Parameter	Observed Power ^a
					Lower Bound	Upper Bound			
Intercept ^{vi}	51.767	4.327	11.963	.000	43.263	60.271	.242	11.963	1.000
RiskScore2	3.986	1.111	3.589	.000	1.803	6.168	.028	3.589	.947
GovTrust	-1.872	.569	-3.288	.001	-2.991	-.753	.024	3.288	.907
MedTrust	-1.895	.620	-3.056	.002	-3.114	-.676	.020	3.056	.862
[UNCTYTotal=000]	-7.032	4.441	-1.583	.114	-	1.696	.006	1.583	.352
[UNCTYTotal=001]	-1.118	5.440	-.206	.837	-	9.573	.000	.206	.055
[UNCTYTotal=010]	-4.382	5.440	-.805	.421	-	6.310	.001	.805	.127
[UNCTYTotal=100]	-2.182	4.441	-.491	.623	-	6.546	.001	.491	.078
[UNCTYTotal=101]	-4.051	4.440	-.912	.362	-	4.674	.002	.912	.149

^{vi} The B coefficient of the Intercept represents the constant in the linear equation.

[UNCTYTotal=110]	-2.881	4.443	-.648	.517	-	5.850	.001	.648	.099
[UNCTYTotal=111]	0 ^b	.	.	.	11.612
[Malignancy=1]	6.700	4.439	1.509	.132	-2.024	15.424	.005	1.509	.325
[Malignancy=2]	0 ^b
[Order=A]	-7.774	4.451	-1.747	.081	-	.973	.007	1.747	.414
[Order=B]	0 ^b	.	.	.	16.520
[UNCTYTotal=000] *	7.700	6.278	1.227	.221	-4.638	20.038	.003	1.227	.232
[Malignancy=1]	0 ^b
[UNCTYTotal=000] *	0 ^b
[Malignancy=2]	0 ^b
[UNCTYTotal=001] *	10.100	7.689	1.314	.190	-5.011	25.211	.004	1.314	.259
[Malignancy=1]	0 ^b
[UNCTYTotal=001] *	0 ^b
[Malignancy=2]	0 ^b
[UNCTYTotal=010] *	9.900	7.689	1.288	.199	-5.211	25.011	.004	1.288	.250
[Malignancy=1]	0 ^b
[UNCTYTotal=010] *	0 ^b
[Malignancy=2]	0 ^b
[UNCTYTotal=100] *	5.700	6.278	.908	.364	-6.638	18.038	.002	.908	.148
[Malignancy=1]	0 ^b
[UNCTYTotal=100] *	0 ^b
[Malignancy=2]	0 ^b
[UNCTYTotal=101] *	1.100	6.278	.175	.861	-	13.438	.000	.175	.054
[Malignancy=1]	0 ^b	.	.	.	11.238
[UNCTYTotal=101] *	0 ^b
[Malignancy=2]	0 ^b
[UNCTYTotal=110] *	6.000	6.278	.956	.340	-6.338	18.338	.002	.956	.159
[Malignancy=1]	0 ^b
[UNCTYTotal=110] *	0 ^b
[Malignancy=2]	0 ^b
[UNCTYTotal=111] *	0 ^b
[Malignancy=1]	0 ^b
[UNCTYTotal=111] *	0 ^b
[Malignancy=2]	0 ^b
[UNCTYTotal=000] *	3.861	6.297	.613	.540	-8.514	16.236	.001	.613	.094
[Order=A]	0 ^b
[UNCTYTotal=000] *	0 ^b
[Order=B]	0 ^b
[UNCTYTotal=001] *	-2.494	7.697	-.324	.746	-	12.633	.000	.324	.062
[Order=A]	0 ^b	.	.	.	17.621
[UNCTYTotal=001] *	0 ^b
[Order=B]	0 ^b
[UNCTYTotal=010] *	4.094	7.697	.532	.595	-	19.221	.001	.532	.083
[Order=A]	0 ^b	.	.	.	11.033
[UNCTYTotal=010] *	0 ^b
[Order=B]	0 ^b
[UNCTYTotal=100] *	.261	6.297	.041	.967	-	12.636	.000	.041	.050
[Order=A]	0 ^b	.	.	.	12.114
[UNCTYTotal=100] *	0 ^b
[Order=B]	0 ^b
[UNCTYTotal=101] *	6.529	6.296	1.037	.300	-5.844	18.903	.002	1.037	.179
[Order=A]	0 ^b
[UNCTYTotal=101] *	0 ^b
[Order=B]	0 ^b
[UNCTYTotal=110] *	-1.018	6.285	-.162	.871	-	11.333	.000	.162	.053
[Order=A]	0 ^b	.	.	.	13.369

[UNCTYTotal=110] * [Order=B]	0 ^b
[UNCTYTotal=111] * [Order=A]	0 ^b
[UNCTYTotal=111] * [Order=B]	0 ^b
[Malignancy=1] * [Order=A]	2.100	6.278	.335	.738	-	14.438	.000	.335	.063
[Malignancy=1] * [Order=B]	0 ^b	.	.	.	10.238
[Malignancy=2] * [Order=A]	0 ^b
[Malignancy=2] * [Order=B]	0 ^b
[UNCTYTotal=000] * [Malignancy=1] * [Order=A]	4.250	8.878	.479	.632	-	21.698	.001	.479	.077
[UNCTYTotal=000] * [Malignancy=1] * [Order=B]	0 ^b	.	.	.	13.198
[UNCTYTotal=000] * [Malignancy=2] * [Order=A]	0 ^b
[UNCTYTotal=000] * [Malignancy=2] * [Order=B]	0 ^b
[UNCTYTotal=001] * [Malignancy=1] * [Order=A]	-1.400	10.874	-.129	.898	-	19.970	.000	.129	.052
[UNCTYTotal=001] * [Malignancy=1] * [Order=B]	0 ^b	.	.	.	22.770
[UNCTYTotal=001] * [Malignancy=2] * [Order=A]	0 ^b
[UNCTYTotal=001] * [Malignancy=2] * [Order=B]	0 ^b
[UNCTYTotal=010] * [Malignancy=1] * [Order=A]	2.400	10.874	.221	.825	-	23.770	.000	.221	.056
[UNCTYTotal=010] * [Malignancy=1] * [Order=B]	0 ^b	.	.	.	18.970
[UNCTYTotal=010] * [Malignancy=2] * [Order=A]	0 ^b
[UNCTYTotal=010] * [Malignancy=2] * [Order=B]	0 ^b
[UNCTYTotal=100] * [Malignancy=1] * [Order=A]	4.500	8.878	.507	.613	-	21.948	.001	.507	.080
[UNCTYTotal=100] * [Malignancy=1] * [Order=B]	0 ^b	.	.	.	12.948
[UNCTYTotal=100] * [Malignancy=2] * [Order=A]	0 ^b

[UNCTYTotal=100] * [Malignancy=2] * [Order=B]	0 ^b
[UNCTYTotal=101] * [Malignancy=1] * [Order=A]	1.600	8.878	.180	.857	-15.848	19.048	.000	.180	.054
[UNCTYTotal=101] * [Malignancy=1] * [Order=B]	0 ^b
[UNCTYTotal=101] * [Malignancy=2] * [Order=A]	0 ^b
[UNCTYTotal=101] * [Malignancy=2] * [Order=B]	0 ^b
[UNCTYTotal=110] * [Malignancy=1] * [Order=A]	1.000	8.878	.113	.910	-16.448	18.448	.000	.113	.051
[UNCTYTotal=110] * [Malignancy=1] * [Order=B]	0 ^b
[UNCTYTotal=110] * [Malignancy=2] * [Order=A]	0 ^b
[UNCTYTotal=110] * [Malignancy=2] * [Order=B]	0 ^b
[UNCTYTotal=111] * [Malignancy=1] * [Order=A]	0 ^b
[UNCTYTotal=111] * [Malignancy=1] * [Order=B]	0 ^b
[UNCTYTotal=111] * [Malignancy=2] * [Order=A]	0 ^b
[UNCTYTotal=111] * [Malignancy=2] * [Order=B]	0 ^b

a. Computed using alpha = .05

b. This parameter is set to zero because it is redundant.

Uncertainty (not significant):

Estimates

Dependent Variable: Total Investment

Uncertainty	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
Baseline (K)	48.674 ^a	1.572	45.585	51.763
Noise (K)	51.198 ^a	2.234	46.809	55.588
Delay (K)	52.078 ^a	2.242	47.672	56.485
Baseline (U)	50.787 ^a	1.572	47.698	53.875
Noise (U)	49.027 ^a	1.579	45.924	52.129
Delay (U)	48.723 ^a	1.579	45.621	51.826
Delay&Noise (U)	48.863 ^a	1.572	45.775	51.952

a. Covariates appearing in the model are evaluated at the following values: Risk Score 2 = 1.9823, GovTrust = 3.48, MedTrust = 2.26.

Pairwise Comparisons

Dependent Variable: Total Investment

(I) Uncertainty	(J) Uncertainty	Mean Difference (I-J)	Std. Error	Sig. ^a	95% Confidence Interval for Difference ^a	
					Lower Bound	Upper Bound
Baseline (K)	Noise (K)	-2.524	2.730	.356	-7.890	2.841
	Delay (K)	-3.404	2.744	.215	-8.797	1.989
	Baseline (U)	-2.113	2.220	.342	-6.475	2.250
	Noise (U)	-.352	2.225	.874	-4.724	4.019
	Delay (U)	-.049	2.231	.982	-4.433	4.335
	Delay&Noise (U)	-.189	2.226	.932	-4.563	4.185
Noise (K)	Baseline (K)	2.524	2.730	.356	-2.841	7.890
	Delay (K)	-.880	3.186	.783	-7.141	5.382
	Baseline (U)	.412	2.730	.880	-4.954	5.777
	Noise (U)	2.172	2.721	.425	-3.176	7.519
	Delay (U)	2.475	2.749	.368	-2.928	7.878
	Delay&Noise (U)	2.335	2.732	.393	-3.034	7.704
Delay (K)	Baseline (K)	3.404	2.744	.215	-1.989	8.797
	Noise (K)	.880	3.186	.783	-5.382	7.141
	Baseline (U)	1.292	2.744	.638	-4.101	6.684
	Noise (U)	3.052	2.761	.270	-2.375	8.478
	Delay (U)	3.355	2.723	.219	-1.996	8.706
	Delay&Noise (U)	3.215	2.732	.240	-2.154	8.584
Baseline (U)	Baseline (K)	2.113	2.220	.342	-2.250	6.475
	Noise (K)	-.412	2.730	.880	-5.777	4.954
	Delay (K)	-1.292	2.744	.638	-6.684	4.101
	Noise (U)	1.760	2.225	.429	-2.612	6.132
	Delay (U)	2.063	2.231	.356	-2.321	6.448
	Delay&Noise (U)	1.923	2.226	.388	-2.451	6.298
Noise (U)	Baseline (K)	.352	2.225	.874	-4.019	4.724
	Noise (K)	-2.172	2.721	.425	-7.519	3.176
	Delay (K)	-3.052	2.761	.270	-8.478	2.375
	Baseline (U)	-1.760	2.225	.429	-6.132	2.612
	Delay (U)	.303	2.246	.893	-4.110	4.717
	Delay&Noise (U)	.163	2.231	.942	-4.221	4.548
Delay (U)	Baseline (K)	.049	2.231	.982	-4.335	4.433
	Noise (K)	-2.475	2.749	.368	-7.878	2.928
	Delay (K)	-3.355	2.723	.219	-8.706	1.996
	Baseline (U)	-2.063	2.231	.356	-6.448	2.321
	Noise (U)	-.303	2.246	.893	-4.717	4.110
	Delay&Noise (U)	-.140	2.225	.950	-4.512	4.232
Delay&Noise (U)	Baseline (K)	.189	2.226	.932	-4.185	4.563
	Noise (K)	-2.335	2.732	.393	-7.704	3.034
	Delay (K)	-3.215	2.732	.240	-8.584	2.154
	Baseline (U)	-1.923	2.226	.388	-6.298	2.451
	Noise (U)	-.163	2.231	.942	-4.548	4.221
	Delay (U)	.140	2.225	.950	-4.232	4.512

Based on estimated marginal means

Pairwise Comparisons

Dependent Variable: Total Investment

(I) Uncertainty	(J) Uncertainty	Mean Difference (I-J)	Std. Error	Sig. ^a	95% Confidence Interval for Difference ^a	
					Lower Bound	Upper Bound
Baseline (K)	Noise (K)	-2.524	2.730	.356	-7.890	2.841
	Delay (K)	-3.404	2.744	.215	-8.797	1.989
	Baseline (U)	-2.113	2.220	.342	-6.475	2.250
	Noise (U)	-.352	2.225	.874	-4.724	4.019
	Delay (U)	-.049	2.231	.982	-4.433	4.335
	Delay&Noise (U)	-.189	2.226	.932	-4.563	4.185
Noise (K)	Baseline (K)	2.524	2.730	.356	-2.841	7.890
	Delay (K)	-.880	3.186	.783	-7.141	5.382
	Baseline (U)	.412	2.730	.880	-4.954	5.777
	Noise (U)	2.172	2.721	.425	-3.176	7.519
	Delay (U)	2.475	2.749	.368	-2.928	7.878
	Delay&Noise (U)	2.335	2.732	.393	-3.034	7.704
Delay (K)	Baseline (K)	3.404	2.744	.215	-1.989	8.797
	Noise (K)	.880	3.186	.783	-5.382	7.141
	Baseline (U)	1.292	2.744	.638	-4.101	6.684
	Noise (U)	3.052	2.761	.270	-2.375	8.478
	Delay (U)	3.355	2.723	.219	-1.996	8.706
	Delay&Noise (U)	3.215	2.732	.240	-2.154	8.584
Baseline (U)	Baseline (K)	2.113	2.220	.342	-2.250	6.475
	Noise (K)	-.412	2.730	.880	-5.777	4.954
	Delay (K)	-1.292	2.744	.638	-6.684	4.101
	Noise (U)	1.760	2.225	.429	-2.612	6.132
	Delay (U)	2.063	2.231	.356	-2.321	6.448
	Delay&Noise (U)	1.923	2.226	.388	-2.451	6.298
Noise (U)	Baseline (K)	.352	2.225	.874	-4.019	4.724
	Noise (K)	-2.172	2.721	.425	-7.519	3.176
	Delay (K)	-3.052	2.761	.270	-8.478	2.375
	Baseline (U)	-1.760	2.225	.429	-6.132	2.612
	Delay (U)	.303	2.246	.893	-4.110	4.717
	Delay&Noise (U)	.163	2.231	.942	-4.221	4.548
Delay (U)	Baseline (K)	.049	2.231	.982	-4.335	4.433
	Noise (K)	-2.475	2.749	.368	-7.878	2.928
	Delay (K)	-3.355	2.723	.219	-8.706	1.996
	Baseline (U)	-2.063	2.231	.356	-6.448	2.321
	Noise (U)	-.303	2.246	.893	-4.717	4.110
	Delay&Noise (U)	-.140	2.225	.950	-4.512	4.232
Delay&Noise (U)	Baseline (K)	.189	2.226	.932	-4.185	4.563
	Noise (K)	-2.335	2.732	.393	-7.704	3.034
	Delay (K)	-3.215	2.732	.240	-8.584	2.154
	Baseline (U)	-1.923	2.226	.388	-6.298	2.451
	Noise (U)	-.163	2.231	.942	-4.548	4.221
	Delay (U)	.140	2.225	.950	-4.232	4.512

Based on estimated marginal means

a. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

Malignancy (significant):

Estimates

Dependent Variable: Total Investment

Function shape	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
Linear	57.116 ^a	.951	55.247	58.986
Non-linear	42.698 ^a	.951	40.829	44.568

a. Covariates appearing in the model are evaluated at the following values:
Risk Score 2 = 1.9823, GovTrust = 3.48, MedTrust = 2.26.

Pairwise Comparisons

Dependent Variable: Total Investment

(I) Function shape	(J) Function shape	Mean Difference (I-J)	Std. Error	Sig. ^a	95% Confidence Interval for Difference ^a	
					Lower Bound	Upper Bound
Linear	Non-linear	14.418 [*]	1.345	.000	11.774	17.062
Non-linear	Linear	-14.418 [*]	1.345	.000	-17.062	-11.774

Based on estimated marginal means

*. The mean difference is significant at the .05 level.

a. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

Order (significant):

Estimates

Dependent Variable: Total Investment

Order	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
A	47.789 ^a	.954	45.913	49.665
B	52.026 ^a	.954	50.150	53.901

a. Covariates appearing in the model are evaluated at the following values: Risk Score 2 = 1.9823, GovTrust = 3.48, MedTrust = 2.26.

Pairwise Comparisons

Dependent Variable: Total Investment

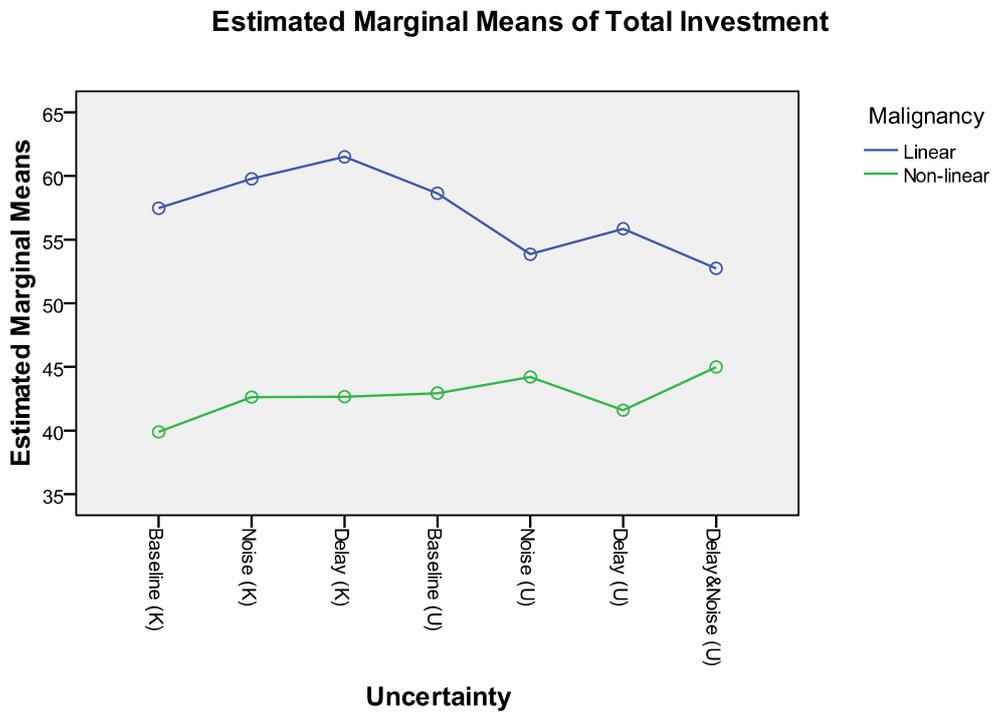
(I) Order	(J) Order	Mean Difference (I-J)	Std. Error	Sig. ^a	95% Confidence Interval for Difference ^a	
					Lower Bound	Upper Bound
A	B	-4.237 [*]	1.354	.002	-6.897	-1.576
B	A	4.237 [*]	1.354	.002	1.576	6.897

Based on estimated marginal means

*. The mean difference is significant at the .05 level.

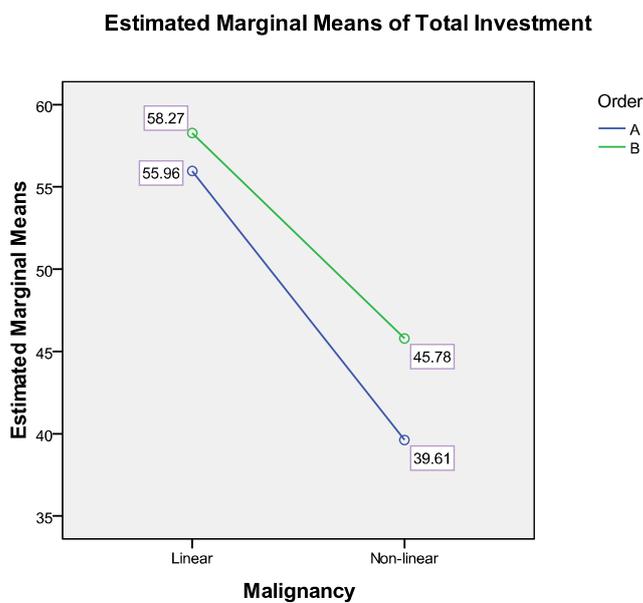
a. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

Uncertainty * Malignancy (not significant):



Covariates appearing in the model are evaluated at the following values: Risk Score 2 = 1.9823, GovTrust = 3.48, MedTrust = 2.26

Order * Malignancy (not significant):



Covariates appearing in the model are evaluated at the following values: Risk Score 2 = 1.9823, GovTrust = 3.48, MedTrust = 2.26

7) Curve Fit: Risk Score 2->Final Damage Level, File Split by Malignancy

Model Description

Model Name	MOD_1
Dependent Variable	1 Final Damage Level
Equation	1 Linear
Independent Variable	Risk Score 2
Constant	Included
Variable Whose Values Label Observations in Plots	Unspecified

Case Processing Summary

Function shape	N	
Linear	Total Cases	240
	Excluded Cases ^a	0
	Forecasted Cases	0
	Newly Created Cases	0
Non-linear	Total Cases	240
	Excluded Cases ^a	0
	Forecasted Cases	0
	Newly Created Cases	0

a. Cases with a missing value in any variable are excluded from the analysis.

Model Summary

Function shape	R	R Square	Adjusted R Square	Std. Error of the Estimate
Linear	.083	.007	.003	13.248
Non-linear	.206	.042	.038	6.478

The independent variable is Risk Score 2.

ANOVA

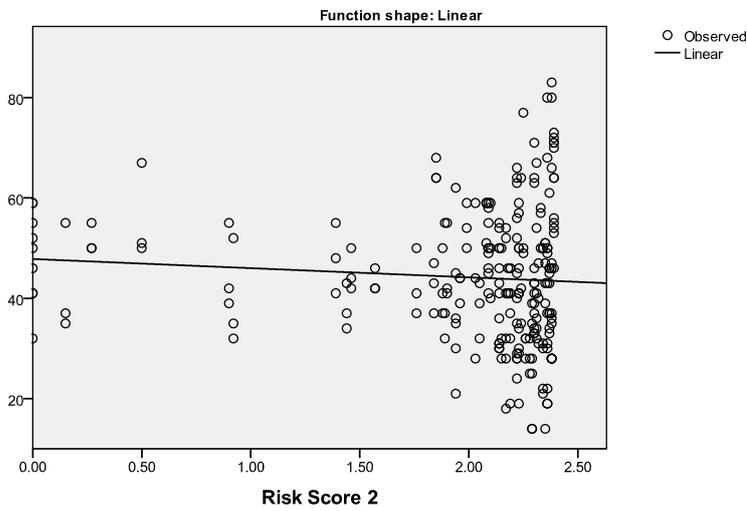
Function shape		Sum of Squares	df	Mean Square	F	Sig.
Linear	Regression	288.746	1	288.746	1.645	.201
	Residual	41773.417	238	175.519		
	Total	42062.162	239			
Non-linear	Regression	442.774	1	442.774	10.550	.001
	Residual	9988.222	238	41.967		
	Total	10430.996	239			

The independent variable is Risk Score 2.

Coefficients

Function shape		Unstandardized Coefficients		Standardized Coefficients	t	Sig.
		B	Std. Error	Beta		
Linear	Risk Score 2	-1.827	1.425	-.083	-1.283	.201
	(Constant)	47.834	2.950		16.213	.000
Non-linear	Risk Score 2	-2.263	.697	-.206	-3.248	.001
	(Constant)	28.481	1.443		19.741	.000

Final Damage Level



Final Damage Level

