Essays on Dynamic Decision Making under Ambiguity

Konstantinos Georgalos
Department of Economics and Related Studies
The University of York

This dissertation is submitted for the degree of
Doctor of Philosophy

October, 2014

Abstract

The first chapter presents an extended survey of the literature on dynamic decision making under ambiguity, focusing both on the theoretical modelling and the available empirical findings. The second chapter, experimentally investigates individual choice under ambiguity in a dynamic setting. Assuming that people have non-Expected Utility preferences, the study is aiming to understand how people update their prior beliefs in a sequential problem. Three different types of decision makers are identified: resolute, naive and sophisticated. In the third chapter, three alternative ways to model stochastic decision making when the choice variable is continuous are presented. These specifications are then tested with data collected from a tailor-made economic experiment.

Contents

List of Figures

List of Tables

Acknowledgements

I would like to express my deep appreciation and gratitude to my supervisor John D. Hey. When I arrived at York in 2011, I only knew John Hey as a distinguished researcher. During the years that I worked with him, I was lucky enough to get to know John also as a person. The term supervisor is quite weak to describe John's merits. I would like to thank him for all the fruitful discussions and constructive controversies we had, for all the time he devoted to ensure that I have everything I needed for my research and for all the invaluable advice that he offered me. I wish that all the future graduate students, could be lucky enough to meet a supervisor with the qualities and characteristics of John. I owe sincere and earnest thankfulness to Jacco Thijssen from the University of York and to Phoebe Koundouri from the Athens University of Economics and Business for their help and support. I would also like to thank Xueqi Dong, Wenting Zhou and Yu-Hsua Lin for their help during the experimental sessions. I am truly indebted and thankful to the Greek Scholarship Foundation that during one of the most difficult times in the modern history of Greece, supported me financially to complete my studies. This support is gratefully recognised. In addition, the financial support from both the Research and Impact Support (RIS) Fund at the University of York and the Department of Economics and Related Studies, for conducting the experiments and participating to conferences, is gratefully recognised.

One of the objectives of this thesis was to make most of the analysis and the writing up using open software (mainly Python, R and LaTeX). Wherever possible, the appropriate credit has been given. Nevertheless, there is a number of anonymous users, who voluntarily provided support and solutions to other software users. To all these people I owe a big "thank you" and I hope one day I will be able to contribute back to this wonderful e-society.

I am grateful to Glenn W. Harisson and Jia Chen for serving on my examination committee.

Finally, I need to thank the people outside academia that without their support nothing would be possible to become real. A *grand merci avec tout mon coeur* belongs to Leslie Pateloup. She experienced every single day of this thesis writing process. She was always there to support me and to encourage me during all good and hard times. Thank you for being there.

Last but not least, I would like to thank my family Antonis, Maria and Mariza for being always there for me. Thank you for always 'nudging' me to the correct path. Thank you for all of your love, support, help, encouragement and dedication.

Declaration

I, Konstantinos Georgalos, declare that this thesis titled, 'Essays on Dynamic Decision Making under Ambiguity' and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made reference to that.

Signed:

____

Date:

____

Konstantinos Georgalos

October, 2014

Introduction

This thesis focuses on issues of *Dynamic Decision Making under Ambiguity.* The term ambiguity is used in the sense that was initiated in the literature by Knight (1921) and Keynes (1921) and describes situations different than risk. In risky choice, a probability distribution exists and it is well-defined and known to the decision maker, while in ambiguous environments there is a lack of a similar distribution. Decision making under ambiguity consists one of the most rapid-growing topics in economic theory and this is not quite surprising, since most of the decisions in actual economic life need to be made in environments where the available information is limited and the decision maker is expected to form some kind of prior beliefs.

The standard model in economic theory, *Subjective Expected Utility* (SEU, Savage (1954)) assumes that the decision makers are Bayesian, which means that they are able to form subjective beliefs in the form of an additive probability distribution and that they make decisions by maximising the Expected Utility based on their beliefs and the outcomes at each possible state of the world. In addition, this model is easily extended to its dynamic framework, as it suffices to assume that these prior beliefs are updated in line with the *Bayesian* updating rule, which ensures the dynamic consistency in choice, a standard assumption in both the microeconomics and macroeconomics literature. Ellsberg (1961) with his famous paradox, challenged the validity of SEU and the result was a vast literature of non-Expected Utility models with the respective experimental studies to emerge. A direct extension of these models is their dynamic version and how they can cope with information reception and updating of prior beliefs. Although the theoretical contributions on this topic are quite extended, there is a lack of empirical studies to confirm or reject the various models and updating rules that have been suggested.

The scope of this thesis is to provide some insights on how people update their prior beliefs and how they decide in the presence of ambiguity.

This thesis consists of three chapters. The first chapter provides a comprehensive literature review on dynamic choice when ambiguity characterises the future states of the world. The main scope is to present a coherent survey on the problems of consistency of behaviour in ambiguous environments, and critically discuss the different approaches that have been suggested in order to interpret these inconsistencies. The survey is then complemented with a discussion on the empirical findings that are available on this issue and a presentation of the updating rules for one of the most eminent family of models of decision making under ambiguity (MaxMin Expected Utility).

In the second chapter, building on the conclusions of chapter 1, that when subjects exhibit attitudes towards ambiguity (either aversion or preference for ambiguity) then preferences cannot be represented by the standard Expected Utility model, we present the design and the results of an economic experiment on dynamic decision making under ambiguity. In this experiment we aim to understand how people behave in a dynamic problem under ambiguity, how do they update their prior beliefs and which is the alternative, non-EU model that best captures all these issues. One of the novelties of this experiment is that it deviates from the standard way that ambiguity is represented in the lab and instead of using standard Ellsberg type urns, a Bingo Blower was used. This device is a transparent, non-manipulable way to represent ambiguity, which eliminates the suspicion that the Ellsberg-type urns generate. Then, we ask a series of allocation problems to the subjects in a sequential choice problem. Based on the data gathered from the experiment, we specify different types of decision makers and we fit the data to different preference functionals and updating rules. We identify three types of decision makers, the resolute, the naive and the sophisticated.

In the third chapter, we build on a methodological problem that was created in the analysis of the experimental data in chapter 2 in the modelling of stochastic choice. As the decision task includes allocation problems (instead of pairwise choices, Holt-Laury price lists, Holt and Laury (2002) and the Becker-DeGroot-

Marschak mechanism, Becker et al. (1964)) with two of more assets, standard ways of modelling stochastic choice are not any more applicable. When the optimal allocation is zero or negative, and a CARA utility function is assumed, then technical difficulties render the estimation of the assumed preference functional impossible as the underlying distributions degenerate. In addition, assuming a power function to represent utility, is a quite restricting assumption as it rules out behavioural patterns (boundary allocations)that are reasonable to be expected during an experimental session or even in real-life economic applications. In this chapter, we present two alternatives to the CRRA modelling specifications for the stochastic term, that allow for boundary allocations. We run an extended simulation and an economic experiment to compare the three proposed stochastic specifications and to understand the consequences that mis-specification in the stochastic choice modelling has.

Chapter 1

A Survey on Dynamic Decision Making under Ambiguity: Theoretical Findings and Experimental Evidence

1.1 Introduction

The scope of this chapter is to provide a comprehensive literature review on dynamic choice when ambiguity for the future states of the world is present. It is taken for granted that the issues that this survey aims to cover, require space that considerably exceeds the size of a doctoral thesis. Nevertheless the main scope is to present a coherent survey on the problems of consistency of behaviour in ambiguous environments, and critically discuss the different approaches that have been suggested in order to interpret these inconsistencies. The survey is then complemented with a discussion on the empirical findings that are available on this issue and a presentation of the updating rules for one of the most eminent family of models of decision making under ambiguity (MaxMin Expected Utility). Similar work has been done by Al-Najjar and Weistein (2009) and to a less extended degree by Klibanoff and Hanany (2007) and Siniscalchi (2011) and more recently by Hammond and Zank (2014). In all the above works, the discussion is constrained to comparing the different ways that deviations from Subjective Expected

Utility (SEU) in a dynamic problem can be accommodated. Nonetheless, they do not proceed to discuss the behavioural implications of certain approaches, neither the descriptive validity they have. The chapter is organised as follows: In the next section, the problem of decision making under ambiguity in a static framework is presented. The aim is to introduce some basic notions, upon which the topic of the thesis is based. Then, the standard problem is extended to its dynamic, or differently, its sequential version and the subsequent behavioural anomalies that emerge are explained. In section 1.3, the different theoretical approaches to overcome similar anomalies are illustrated. In section 1.4, we present the different types of decision makers that can be identified, depending on the way they cope with the dynamic problem. Section 1.5 presents the main empirical findings based on experimental results. The experiments are divided in four categories, dynamic choice under risk, experiments that aim to classify different types of decision makers, dynamic choice under ambiguity and finally experiments that test learning under ambiguity. In section 1.6, we present an example of how different updating rules are applied to the 3-colour Ellsberg paradox, when MaxMin Expected Utility preferences are assumed. Then we conclude.

1.2 Theoretical Framework and Notation

Before starting, it is necessary to introduce some notation and some definitions that will be useful for the presentation of the axioms to follow. We consider a set S as the set of states of the world and a Σ-algebra of subsets of S and denote X as the set of consequences. Let $\mathcal{F}$ the set of all the simple acts, all the Σ-measurable functions $f : S \to X$. X is defined as the subset of constant acts in $\mathcal{F}$ such that, for any $x \in X$, $x(s) = s$ for any $s \in S$. For any $f, g \in \mathcal{F}$ and $A \in \Sigma$, fAg denotes the act which yields $f(s)$ for $s \in A$ and $g(s)$ for $s \in A^c \equiv S \backslash A$. Also we define the preference relationship $\succsim$ and we for the moment we assume a von Neumann-Morgenstern utility function $u : X \to R$.

A prospect is denoted as $(E_1, x_1; E_2, x_2, \ldots; E_n x_n)$ where the outcome x_s is obtained if the event E_s happens. The idea of ambiguity is that there is a series of events $E_1, E_2, \ldots E_n$ and only one will happen but this is not known with certainty and there is a lack of an objective probability distribution over the different events.

The standard assumptions of weak ordering hold in this framework. The decision maker is endowed with a preference relationship $\succsim$ over the different prospects which is complete and transitive. Similarly, $\succ$ defines strict preference and $\sim$ defines indifference. The certainty equivalent is defined as the sure amount of income that makes the decision maker to be indifferent between this amount and the prospect.

Then, since the focus is on dynamic or sequential choices that are realised after the acquisition of some relevant information, the preference relationship should be extended to its conditional form. Thus, by $\succsim_E$ we refer to a preference relation on $\mathcal{F}$, which represents the preferences of the decision maker based on the updated beliefs after receiving the information related to E.

1.2.1 The Ellsberg Paradox

The thought experiment that Daniel Ellsberg presented in his seminal paper on 1961 became the departure point for a vast literature that challenges the famous *Savage Axiom*, or the *Sure-Thing Principle*, but also the notion of *probabilistic sophistication* (both to be explained later). The main message of these experiments is that regardless of the way the subjects decide, it is impossible to assume that they act according to a well-defined probability distribution. The idea was that due to the presence of ambiguity aversion the subjects experience preference reversals. The decision task is the following:

Ellsberg's Two-colour Urn

In the first experiment (*Two-Colour Ellsberg Paradox*) the subject is asked to choose between bets that involve two urns, urn I and urn II. Each urn has 100 balls. In urn I it is known that 50 balls are Black and 50 are Red. In urn II there is no information given concerning its composition. The subject needs to decide between the bets that are shown on Table 1.1.

Bet f pays 100 monetary units if a Red ball from urn 1 is drawn (denote this as R_1) and zero otherwise. Similarly, bet g offers a prize of 100 if a Red ball is drawn from urn 2 and zero otherwise (R_2). f', g' are symmetrically defined as (B_1, B_2). Then, a subject is asked to choose between the following bets: (R_1, B_1), (R_2, B_2), (R_1, R_2) and (B_1, B_2). Focusing on the last two bets, a common response to

Table 1.1 Two-colour Ellsberg Urn

	URN I		URN II	
bet	r	b	r	b
f	100	0		
g			100	0
f′	0	100		
g′			0	100

Table 1.2 Three-colour Ellsberg Urn

	30	60	
bet	b	r	y
f	100	0	0
g	0	100	0
f′	100	0	100
g′	0	100	100

these bets is that: $(R_1 \succ R_2)$ and $(B_1 \succ B_2)$. Choosing so, means a direct violation of the axioms of Expected Utility (EU). In this experiment if we obtain that $g \prec f$ and $g' \prec f'$ this is a violation of the requirement that the probabilities should sum up to 1 since if the preferences are like above, it holds that $P(R_1) < P(R_2)$ and$P(B_1) < P(B_2)$ violating the relationship that

$$P(R_1) + P(B_1) = P(R_2) + P(B_2) = 1$$

This pattern is inconsistent with the *Sure-Thing Principle* leading to contradiction with the subjective Expected Utility model, but it is also a violation of the *probabilistic sophistication* as it is defined by Machina and Schmeidler (1992).

Ellsberg's Three-colour Urn

In the second experiment (*Three-Color Ellsberg Paradox*), urn I has 90 balls of which 30 are Black (B) and the rest 60 are Red (R) and Yellow (Y) (unknown proportions) and could be any number between 0 and 60. The decision maker (DM) has to decide between the following bets that pay the amounts below:

The DM firstly has to decide between the bets f and g and then between f' and g'. Empirically, subjects prefer f to g and g' to f' which is a violation of the Sure-Thing Principle since the bets f' and g' are obtained by changing the common outcome in yellow from 0 to 100. Using the SEU and assigning to a normalised

utility function such that $u(100) = 1$ and $u(0) = 0$ it holds that:
if $f \succ g$ then $P_b > P_r$ and if $g' \succ f'$ then $P_r + P_y > P_b + P_y$ which generates a contradiction that leads to decisions incompatible with SEU. Following Hammond and Zank (2014), this can be shown by assuming that there are b Black balls inside the urn. Then, the composition of the urn can be represented by the following vector $(R, B, Y) = (30, b, 60 - b)$ and the respective probabilities attached to the events are $(1/3, p, q)$. Assuming that for each b in $\{1, 2, \cdots, 60\}$ the decision maker forms some kind of subjective beliefs (P_b) that b Black balls are inside the urn. The probability that a Black ball is randomly drawn is given by $p = \sum_{b=0}^{60} p_b b/90$ and similarly for a Yellow ball $p = \sum_{b=0}^{60} p_b(60 - b)/90$. Preferring f to g implies that $\frac{1}{3} > p$ while preferring g' to f' implies that $p + q > \frac{1}{3} + q$ which contradicts with the previous result and ensures the violation of the probabilistic sophistication.

This is a proof of why we cannot expect that the agents are probability sophisticated. The above result is aimed to be explained as a rational choice that violates one of the axioms[1] of the Subjective Expected Utility model due to the existence of attitudes towards ambiguous events. More specifically, speaking in terms of the Savage setup (Savage (1954)), the axiom of the *Sure-Thing Principle*, is violated. This axiom resembles the Independence axiom, when Expected Utility is set up to accommodate risky prospects[2]. Formally, the axiom require that for all acts f, g, h, h' and for every event E,

$$fEh \succsim gEh \Leftrightarrow fEh' \succsim gEh'$$

Verbally, this axiom requires that the preference over two acts f, g should only depend on the values of f, g when they differ. Let for instance, f, g to differ on

[1]The Savage Expected Utility model, requires that a series of rationality axioms are being satisfied. These axioms include, the Weak Ordering of Acts (P_1), the Sure Thing Principle (P_2), the Existence of Intrinsic Preferences (P_3), the Non-Influence of the Prize (P_4), the Non-Triviality of Preferences under Certainty (P_5), the Continuity (P_6), the Dominance (P_7), the Measurability (P_8) and the Event-wise Continuity axiom. When all the axioms above are satisfied, then a von Neumann-Morgenstern utility function exists and preferences can be represented by the Expected Utility model. For an analytical discussion of the axioms see Gilboa (2009).

[2]Camerer and Weber (1992), Camerer (1995) and more recently Starmer (2000), although a bit dated, survey the literature on the experimental evidence of individual choices and the violations of rational behaviour that are observed. A recent review of the literature on this topic is provided by Hey (2014).

an event A. Then, if A does not occur, f, g result in exactly the same outcomes. In other words, when the two acts are compared, it suffices to focus on A and ignore the complement event A^c. What needs to hold is that $f(s) = g(s)$ when $s \notin A$. As a consequence, a vast literature has emerged on decision theory under ambiguity when the problem under investigation is static. Several theories have been developed in order to accommodate Ellsberg-type behaviour, with the most prominent being the Multiple Prior preferences (MaxMin Expected Utility, Gilboa and Schmeidler (1989)), the Choquet Expected Utility (Schmeidler (1989)), the Cumulative Prospect Theory (Tversky and Kahneman (1992)), the α-MaxMin Model (Ghirardato et al. (2004)), the Smooth Preferences model (Klibanoff et al. (2005)), the Vector Expected Utility (Siniscalchi (2009)), the Contraction model (Gajdos et al. (2008)), the Variational Preferences model (Maccheroni et al. (2006a)), to name but a few. For an extended review of the models of decision making under ambiguity see Etner et al. (2012). All the theories mentioned above succeed to accommodate preferences that are able to explain the Ellsberg paradox due to the relaxation of the Sure-Thing Principle. Besides, all these alternative approaches, assume some kind of preference functional which also account for attitudes towards ambiguity. The latter allows for the empirical testing of these theories.

Several experiments have been conducted in order to compare the different models regarding their fitting on the heterogeneous behaviour of the subjects under ambiguity and their predictive power. These studies include a variety of different methods to represent ambiguity as well as different types of questions and decision tasks that allow for the elicitation of beliefs and ambiguity attitudes. Halevy (2007), tests four basic models using reservation values found applying the *Becker-DeGroot-Marschak* mechanism (Becker et al. (1964)) for four different urns without being able to show which model performs best as he did not proceeded econometrically in a way that allows similar inferences. What he did was to perform tests individually on each model to conclude that the Ellsberg paradox is created due to the inability to reduce compound lotteries, a fact that the current models are not able to capture. Andersen et al. (2009), use a task that allows to simultaneously estimate a two-parameter model that captures attitudes towards risk and ambiguity. They found that ambiguity aversion is quantitatively significant and that attitudes towards risk and ambiguity are significantly different. Hey et al. (2010), perform an extensive comparison among the different models by fitting preference functionals and using a set of pair-wise questions. They found that models that lack some kind of preference functional do not perform well, neither do more sophisticated models such as the Choquet Expected Utility. Hayashi and Wada (2010), find remarkable violations of the α-MaxMin model in an experiment with imprecise information which is represented in the form of a set of possible probability values. In Abdellaoui et al. (2011) ambiguity is created using two different devices, an 8-ball Ellsberg-type urn and bets among natural, ambiguous events (e.g. the weather in a known compared to a less known country). They test only Rank Dependent Expected Utility and applying a set of Holt-Laury price lists. They found considerable heterogeneity in subjects' preferences. They also introduce the *source dependent* approach. Ahn et al. (2014), represented ambiguity with an Ellsberg-type urn where they specified three different assets for three respective states of the world, of which the one would occur with probability equal to 1/3, while there was no information on the likelihood of the two remaining states. The problem task was an allocation problem, transforming the experiment to a portfolio choice one. Instead of fitting different preference functionals of the various specifications, they split the models in two large families, namely the *kinked* and the *smooth*. They found remarkable degrees of heterogeneity on individual preferences and high degrees of ambiguity aversion. Also they found that a specific percentage of the subjects adheres to the α-MaxMin model while the behaviour of another significant proportion agrees with the Smooth model. Finally, Hey and Pace (2014), using a set of allocation type questions and representing ambiguity using a Bingo Blower, they test the descriptive power of a set of different theories of decision making under ambiguity. The authors go one step further and use a part of the data to test the predictive power that these models have. They found that the most sophisticated models do not necessarily perform better than the simpler ones.

It is for sure that there is still long way to go before adopting a suitable model for static settings. Nevertheless, the real important issues in all fields of economic theory involve time, acquisition of information, updating of beliefs and sequential choice. Thus it is of paramount importance to identify which of the available models is the most suitable to describe behaviour in a dynamic, ambiguous environment. Gilboa and Schmeidler (1993) stress the importance of belief updating:

1. The theoretical validity of any model of decision making under uncertainty is quite dubious if it cannot cope successfully with the dynamic aspect.

2. The updating problem is at the heart of statistical theory. In fact, it may be viewed as the problem statistical inference is trying to solve. [. . .]

3. Applications of these models to economic and game theory models require some assumptions on how economic agents change their beliefs over time. The question naturally arises, then: What are reasonable ways to update such beliefs?

4. The theory of artificial intelligence, which in general seems to have much in common with the foundations of economic, decision and game theory, also tries to cope with this problem [. . .]

The next example presents a famous paradox, the *Monty Hall* problem. Although the problem refers to risky choice, where the distribution of the probabilities is known, it is quite interesting to motivate this discussion by presenting this updating anomaly as it shares the element of sequential decision making upon the receipt of information.

1.2.2 The Monty Hall Problem

This famous probability puzzle shows the difficulty that decision makers may experience when receive new information. The *Monty Hall* problem firstly stated by Selvin (1975) and is named after by Monty Hall, the host of the TV game show Let's Make a Deal. In this problem, a contestant faces three different doors (let them be doors A, B and C) where one of them hides a prize of high value (e.g. a car) while the other two something valueless. The contestant chooses a door (let us say A) and then the presenter opens one of the other two doors that hides one of the valueless prizes (let it be door B). Then, the contestant is asked if she wishes to switch the initial choice and choose door C or not. The empirical evidence shows that the majority of the participants turns down the switching option (as there are no available data from the actual show, a number of studies have tried to replicate the problem with experimental subjects, with the most prominent being the one conducted by Friedman (1998) who found that in overall only 28.7% of the subjects do change their initial choice). The puzzle appears in the fact that subjects do not seem to be probabilistically sophisticated and consequently, fail to accordingly update their beliefs. The common argument in favor of sticking to the initial decision, is that since *ex-ante* the prize can be behind either of the three doors with equal probability $P(A) = P(B) = P(C) = 1/3$, now that there are only two doors remaining the probability of the prize being behind the chosen door is equal to $P(A|\{A,C\}) = P(C|\{A,C\}) = 1/2$ thus, there is no reason to change. Nevertheless, reasoning like this fails to anticipate that switching actually improves the probabilities of winning. It can be shown that the two posterior probabilities are not identical by applying the Bayesian rule. Let E be the event that the door opened is door B. Then the posterior probabilities for the two states A and C are given by the following formula:

$$P(A|E) = \frac{P(E|A)P(A)}{P(E)}$$

$$P(C|E) = \frac{P(E|C)P(C)}{P(E)}$$

ex-ante it is known that $P(A) = P(C) = 1/3$. The crucial point for the decision maker is to realise that $P(E|A) \neq P(E|C)$. An assumption that is made here is that the decision on which door Monty Hall opens, is not independent on where the car is hidden. The probability that door B is opened when the car is in A is equal to 1/2. Similarly, the probability that the door B is opened when the car is in door C is equal to 1. Consequently, $P(E|A) < P(E|C)$ which means that $P(A|E) < P(C|E)$ and the contestant increases her probability of winning when she decides to switch. To translate these probabilities to numbers, it suffices to calculate $P(E)$ which is simply $P(E) = P(E|A)P(A) + P(E|B)P(B) + P(E|C)P(C) = 1/2 \times 1/3 + 0 \times 1/3 + 1 \times 1/3 = 1/2$ and substituting to the Bayesian formula gives that $P(A|E) = 1/3, P(C|E) = 2/3$

The problem becomes even more interesting in the case where the candidate has the opportunity to explicitly state her actions *ex-ante* as for example to state a plan that says choose door A, switch if door B is opened, stay otherwise. Having similar information allows to test the fundamental assumption in economic theory that application of the Bayes rule, demands subjects to be dynamically consistent. Let us now move to a sequential problem under ambiguity, where the probability distributions are not any more well-defined. The standard way to do so in the literature is by extending the Ellsberg-type thought experiment to its sequential form task that we present below.

1.2.3 The Sequential Ellsberg Paradox

Extending the Ellsberg's paradox in a dynamic framework[3] is useful in order to understand how dynamic consistency is violated. In this setup, the decision maker is assumed to have *ex-ante* and *ex-post* preferences over acts. Consider the Ellsberg three-colour case. In the urn there are 90 balls, 30 of those are Red (R) and the rest Black (B) or Yellow (Y). Again there are three states of the world in the state space $\Omega = \{B, R, Y\}$ and we assume a three-period model where at $t = 0$ the decision maker forms the prior probabilities for each colour (always knowing that the proportion of Red balls is equal to $1/3$), at $t = 1$ she receives the filtration $\mathcal{F}_1 = \{\{R, B\}, \{Y\}\}$ which provides the information whether the ball is Yellow or not and at $t = 2$ all the ambiguity is resolved and the colour of the ball is revealed (in case that the ball was not Yellow). The different bets presented above (Table 1.2), can now be represented as a dynamic problem. In the first case the decision maker must choose between f and g. So writing the payoffs in the form (B, R, Y) she has to decide between $(100,0,0)$ and $(0,100,0)$. Assuming that the decision maker at $t = 0$ typically expresses the following preference:

$$(100,0,0) \succ_0 (0,100,0)$$

Using the standard tree representation firstly presented by Raiffa (1968) where squares denote decision nodes and circles chance nodes the choice between f and g is:

[3] A similar approach can be found in the literature in Epstein and Schneider (2003), Klibanoff and Hanany (2007), Hanany and Klibanoff (2009), and Siniscalchi (2011).

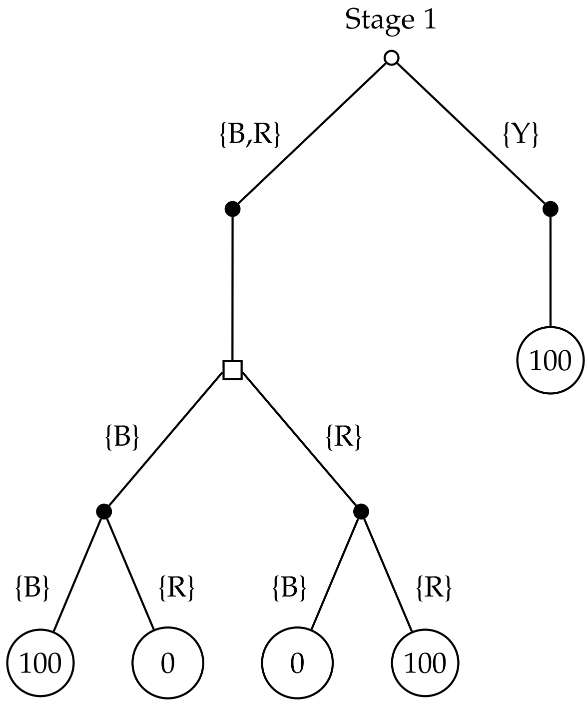

Now, given that a ball is drawn the decision maker has to decide between the bets of getting 100 if the ball is either black (B) or red (R) or getting 100 if the ball is either red (R) or yellow (Y). So writing the payoffs in the form (B, R, Y) she has to decide between $(100, 0, 100)$ and $(0, 100, 100)$. Typically, at $t = 0$ the following preference is expressed:

$$(100, 0, 100) \prec_0 (0, 100, 100)$$

which is the expression of ambiguity aversion as the decision maker seems to prefer the less ambiguous bets. The respective tree in this case is:

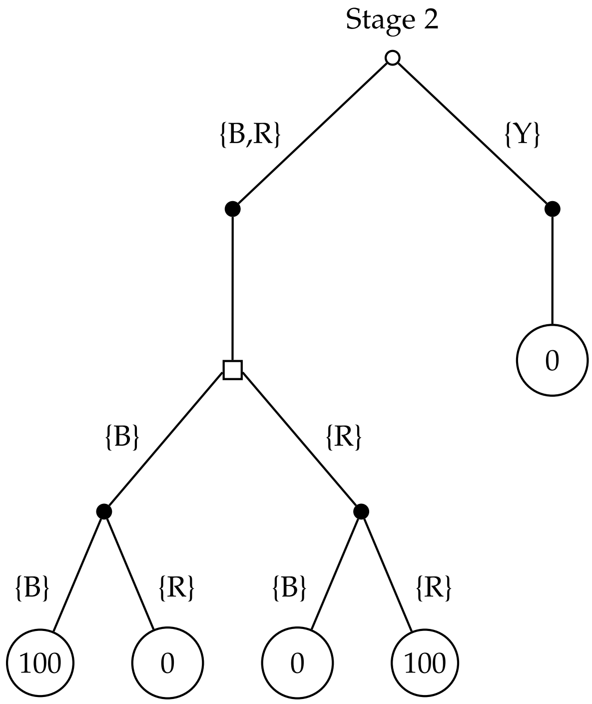

Now if at time $t = 1$ the DM is informed that the ball is not Yellow $\{\neg Y\}$ thus the event $\{R, B\}$ happened, and then she is given the chance to choose between a bet on a Black ball (B) or on a Red ball (R). Following the same line of thinking as before, the preferences that are expressed are:

$$(100, 0, 100) \succ_{1,\{R,B\}} (0, 100, 100)$$

if the ball is not Yellow and

$$(100, 0, 100) \sim_{1,\{Y\}} (0, 100, 100)$$

if the ball is green, which leads to a violation of dynamic consistency and consequently backward induction does not seem to apply anymore. A result which according to the standard theory should not happen since the opportunity to condition the choices on the new information does not change the problem in an essential way (Klibanoff and Hanany (2007)). This leads to the question of how dynamic models of decision making under ambiguity are able to explain similar behaviour, or, in other words, how decision makers update their beliefs in the presence of new information.

Therefore, a crucial question in the theory of decision making under ambiguity is the way that beliefs are updated upon the arrival of new information. In the standard economic theory when a decision maker has to cope with a similar dynamic problem where the various probabilities of the possible states of the world are given (are subjective), it is common to assume that beliefs are updated according to the *Bayesian* rule in the arrival of new information. This is based on the theory of Subjective Expected Utility (SEU, Savage (1954)) where when certain axioms are satisfied, the preferences of the decision maker can be represented by a utility function unique up to a linear positive transformation and her beliefs are represented by a subjective, additive probability measure. In other words, assuming that the decision maker has Expected Utility preferences, this leads to the result that her beliefs are updated in the arrival of new information in a Bayesian way, securing in this way Dynamic Consistency. On the other hand, if the decision maker is dynamically consistent then the preferences are updated in a Bayesian way. This result is stated in the literature by Ghirardato (2002), and it is proven in Klibanoff and Hanany (2007) that Dynamic Consistency is the primary justification for Bayesian updating and under the view that Bayesian updating should be taken as given, dynamic consistency comes "for free" under Expected Utility . In addition, In Epstein and Le Breton (1993), it is shown that when conditional preferences are based on beliefs in a Dynamically Consistent way then the decision maker must be probabilistically sophisticated and has a Bayesian prior which automatically rules out Ellsberg type behaviour.

But before going on, it is necessary to formally define the rationality axioms that should hold, when preferences are characterised by SEU. The axioms of the Savage Expected Utility require among the other standard axioms of weak order preferences, the axiom of *Dynamic Consistency* (DC), the axiom of *Consequentialism* (C) and the axiom of *Reduction of Compound Lotteries (RCL).*[4]

In the literature of dynamic decision making, the axioms of Dynamic Consistency and Consequentialism are also known as rationality criteria in the sense that a decision maker should always comply with the two axioms if she is rational. From a normative point of view, behaviour that is non DC can lead to *dominated* choices. According to Wakker (1998), DC is useful as it ensures that information does not have negative value. DC makes it easier to describe a decision

[4]This holds in decision under risk. When the acts are Savage then a formalisation of the reduction of compound lotteries axiom in a subjective setting is not possible.

maker that plans ahead and according to this we can make welfare statements in dynamic models. From a psychological point of view DC can be viewed as a rationalisation property: DC updating rules are those that support earlier choices or plans. Finally, in general it is assumed that when a decision maker is dynamically consistent then she maximises her welfare[5]. We next define formally the two rationality axioms.

Assume two acts f and g and we want from the prior preferences to predict the conditional preferences when an event E occurs. The axiom of the *Sure-Thing Principle* (P_2), suggests that it suffices to check the unconditional preference between any pair of acts, let them to be f' and g', that agree outside the event E but in the event E, f' agrees with the act f and g' with g. Thus it holds $f' = g'$ if $\omega \in E$ and $f' = f, g' = g$ if $\omega \in E^c$. Then the idea is that only the states in E count for the preferences and consequently, if $f' \succsim g'$ then $f \succsim_E g$[6]. To illustrate this it is useful to consider again the Ellsberg paradox that was presented before. In the Table below there are the payoffs in the event $\{Y\}$. It is easy to see that $f' = g'$ and $f = g$.

bet	y
f	0
g	0
f′	100
g′	100

In a similar manner, in the following Table it can be seen that in the event $\{B, R\}$ $f = f'$ and $g = g'$.

bet	b	r
f	100	0
g	0	100
f′	100	0
g′	0	100

[5]Although Ozdenoren and Peck (2008) using a game theoretical model, show that there are certain situations where it is optimal to be dynamic inconsistent.

[6]This is the analysis in the Savage framework. The same can be presented in the Anscombe-Aumann framework. In this case, instead of Savage acts we have lotteries. The Sure-Thing Principle is defined as: for $f, g, h, k \in L_0$ and $E \in \Sigma, fEh \succsim gEh$, implies that $fEk \succsim gEk$.

Summarising, the idea is that the unconditional preferences are determined only by the differences that appear on the conditioning event since the two actions are equivalent outside this event. This idea is crucial since it incorporates the two rationality criteria that are used in order to explain how the decision makers update their beliefs, dynamic consistency and consequentialism. A DM is dynamically consistent when the *ex-ante* preferences coincide with the *ex-post*. If an act is preferred to another unconditionally and receiving the information that the complementary of an event where the two acts agree has been obtained, the conditional preference should remain the same.

Definition 1. *(Dynamic Consistency)*
For any non-null event E and acts $f, g \in \mathcal{F}$ such that $f(\omega) = g(\omega)$ for each $\omega \in \Omega$, $f \succsim g$ implies $f \succsim_E g$.

The importance of the axiom of Dynamic Consistency can be easily understood if one thinks all the applications in the standard economic theory where time is incorporated in the decision process. The majority of models in macroeconomics assume that the agents or the governments are dynamically consistent and thus an exponential discount factor is applied in order to discount future values. These kind of models includes models in finance, saving models and generally whenever dynamic or long-horizon decisions are analyzed (i.e. environmental economics, retirement plans). In addition, another field of economics where the assumption of DC is crucial, is the dynamic game theoretical models. In the majority of the literature, the players are assumed to be dynamically consistent and to update their beliefs using the standard Bayesian rule.
Accordingly, consequentialism is satisfied when the DM does not take into account states that are not available anymore and thinks of the rest of the decision tree as being a new problem (Hammond (1988)). A DM satisfies consequentialism when the decision maker conditions her preferences on an event E taking into consideration her unconditional preference and treating E^c as a null-event.

Definition 2. *(Consequentialism)*
For any non-null event E and acts $f, g \in \mathcal{F}$ such that $f(\omega) = g(\omega)$ for each $\omega \in \Omega$ implies $f \sim_E g$.

Machina (1989) and McClennen (1990) suggest that consequentialism is to be avoided as this requires some kind of dynamic separability of preferences while ambiguity aversion suggests that preferences do not need to be separable across events. Al-Najjar and Weistein (2009) is using the term "fact-based updating" to describe situations where the DM displays preference reversals and thus experiences a conflict between the *ex-ante* and the *ex-post* preferences. An updated preference is fact-based if for any event E and acts f, g, f', g' if $f \equiv_E f' \& g \equiv_E g'$ then $f \succsim_E g \Leftrightarrow f' \succsim_E g'$. Updating in this way means that when the two acts are compared conditionally on the event E, and no weight is placed on situations outside E or in any other *ex-ante* optimal plan (form of consequentialism). In the example of Ellsberg this can be thought as the updating of the initial set of priors when the DM learns that the event is {R,B} where the conditional preference relation is independent of the acts that are related to {Y}.

The evidence obtained by the Ellsberg paradox shows that the modal preferences violate the Sure-Thing Principle. Therefore, when a decision maker does not have Expected Utility preferences, satisfaction of Bayesian updating, consequentialism and dynamic consistency are not feasible at the same time. As a consequence, a huge literature emerged aiming to model choice that rationally violates the Sure-Thing Principle. In order to be able to do so, one of the rationality axioms DC and C, should be relaxed. In addition, there are theoretical constructions that aspire to preserve both axioms. This comes at a cost regarding the behavioural assumptions that need to be made, as either the feasible interval inside which the priors can be updated, or the information structure that is received by the decision maker. Moreover, each theoretical approach should be accompanied by a story on how beliefs are updated. In the literature, there are three ways that have been suggested, either to preserve dynamic consistency, or to preserve consequentialism, or to impose a recursive structure. We analytically present all the different modelling ways.

1.3 Theoretical Literature

For non-additive beliefs there have been two different approaches of how beliefs are updated. The first one is the *statistical approach* that considers for different updating rules the statistical properties of the updated beliefs that are derived from such rules[7]. The other approach is the *Decision Theoretic* approach where the updating rules arises from axioms on the preferences both conditional and unconditional. We focus on the latter. In this literature, a classification can be done according to the axioms that are preserved in order to explain behaviour. This leads to the inevitable need for some kind of behavioural justification as there are several cases where the choice of the axioms seems to serve the mathematical robustness or some normative objective though it is difficult to explain economic subjects' behaviour. According to this, the different rules that have been suggested in the literature can be categorised in three different approaches: the dynamic inconsistent approaches which preserve consequentialism, the dynamic consistent updating rules (non-consequentialist) and the recursive methods (both dynamic consistent and consequentialist).

1.3.1 Consequentialist

The idea of preserving C and abandoning DC is that the preference conditional on the event E depends only on the unconditional preference, the event E and treats E^c as a null event. Al-Najjar and Weistein (2009) refer to this *naive-updating*, as they show that using similar updating rules may lead to dominated outcomes. All the updating rules that are applied under this framework are dynamically inconsistent. Rules have been proposed for the MaxMin Expected Utility (MEU) by Gilboa and Schmeidler (1993) and for the Choquet Expected Utility model (CEU)[8] by Eichberger et al. (2007) and Eichberger et al. (2010). In section 1.6 an analytical example of the updating rules proposed by Gilboa and Schmeidler (1993) is presented. Pires (2002) axiomatises a Bayesian update rule where all priors are kept, and all are updated according to the Bayes rule. Similar approaches include the work by Wang (2003) and Siniscalchi (2011). A behavioural interpretation of this

[7]For references see Eichberger et al. Eichberger et al. (2007).

[8]These rules are analytically presented in chapter 2.

rule is provided by Eichberger et al. (2007):

> Our reason for retaining consequentialism and dropping dynamically consistency is because, [. . .] we feel ambiguity arises in a fundamental sense from uncertainty about probability created by missing information that is relevant and could be known. Hence once an event is known to have obtained, the only remaining ambiguity the individual faces relates to uncertainty about the probabilities of sub-events of that event. Past (or borne) uncertainty one may have had about the probability of counterfactual event and its subsets are no longer relevant. But such uncertainty might have been relevant to the individual at the time when she did not know whether the event or its complement had obtained, and so such ambiguity that she perceived there to have been *ex-ante*, may well have had an impact on her unconditional preferences.

1.3.2 Dynamic Consistent (non-consequentialist)

These approaches suggest that in order to preserve some form of consistency, the axiom of consequentialism must be dropped. Machina (1989) and McClennen (1990) impose strong requirements for the definition of dynamic consistency, as they require conditional choices to be in agreement with the unconditionally optimal plan. In a similar way, Klibanoff and Hanany (2007), Hanany and Klibanoff (2009) using a weaker version of dynamic consistency, are able to obtain an axiomatisation of dynamic consistent updating rules. In Klibanoff and Hanany (2007), the updating rules are given for the MEU model. Bayes rule is applied to subsets of the priors which subset is a function of the preferences, the conditioning event and the choice problem. In Hanany and Klibanoff (2009), they explore the ways that the models can be extended and try to find the updating rules for more general preferences rather than Expected Utility preferences. They characterise dynamically consistent updating rules for essentially all continuous, monotonic preferences that are ambiguity averse. Using this methodology, they characterise dynamically consistent updating rules satisfying closure for specific models of ambiguity averse preferences. In their approach, they drop consequentialism, without providing any behavioural intuition for doing so and they drop all the problematic priors that could cause reversals satisfying dynamic consistency.

Recursive Methods

This approach was firstly suggested by Epstein and Schneider (2003) and it was applied in the MEU model. The axiomatisation of the preferences preserves both dynamic consistency and consequentialism and this happens due to recursion. In order for recursion to hold, and therefore for the decision maker to be able to use standard dynamic programming techniques, the representing set of measures must be *rectangular*, a quite restrictive assumption that constrains the information structures. In other words, it requires that preferences over acts are recursive and consequently dynamic consistent, in all decision trees consistent with a given filtration. As Al-Najjar and Weistein (2009) describe this approach: "[. . .] to eliminate the updating paradoxes is to limit attention to decision trees (information structures) on which no reversals occur." Consequently, imposing such restrictions to the information filtration rules out Ellsberg type behaviour. Nevertheless, there is no clear behavioural explanation of why one should update the prior beliefs based on this methodology. Similar work using the recursive methodology in several different models of decision making under ambiguity includes the work of Maccheroni et al. (2006b) for the dynamic Variational model, Klibanoff et al. Klibanoff et al. (2009) for the Smooth ambiguity model. Finally, there is the work by Ozdenoren and Peck (2008) who show that in a game theoretical framework and using Ellsberg-type problems, there are cases where the optimal strategy is either to be dynamically inconsistent or consequentialist, depending on the type of the nature that the agent is playing against (i.e. if the nature is perceived as benevolent or malevolent). This can be interpreted as a subgame perfect equilibrium strategy and can rationalise both the violation of dynamic consistency and consequentialism.

1.4 Types

Till now, it has been clear that when decision makers do not have Expected Utility preferences, then they do not update their prior beliefs according to the Bayesian rule and the direct implication of this is the violation of one of the two main axioms, that of dynamic consistency and consequentialism. What seems to be remarkably interesting now, is to identify what people are doing in order to cope with these inconsistencies. All the alternatives that have been presented above, imply either implicitly or explicitly a certain type of behaviour. Al-Najjar and Weistein (2009) discuss the issues on how realistic are the several different rules suggested in the literature from a positive point of view. Recursive methods as well as dynamic consistent updating rules do not seem to apply in real life. As a consequence it is useful to focus on dynamically inconsistent updating rules and to question how this inconsistency is resolved. In this section we describe the three most prominent types[9] that have appeared in the literature and have been exploited in applications such as game theory or information economics to name but a few. The types we consider here are three, the resolute, the naive and the sophisticated.[10] The sophisticated decision maker anticipates that the preferences of the future *ex-post* self are those that are imposed . A sophisticated plan takes into consideration the constraint that choice at the last node must be optimal by the perspective of the conditional preference $\succsim_E$.

1.4.1 Resolute

This type of decision maker solves the problem as if it is a one-period problem. A decision maker who is resolute is supposed not to change her first period decisions and this can be done either violating consequentialism or dynamic consistency. This theory was firstly suggested by Hammond (1988), and then was formalised by McClennen (1990) and Machina (1989) in the context of risk preferences. The idea of a resolute decision maker embraces the notion of *commitment* where it is implicitly assumed that a decision maker realises that violating the contingent plan will have detrimental effects, she chooses to use some kind of commitment device and stick to the *ex-ante* preferences. This is a well known strategy when one has to cope with situations such as temptations or addictions and needs

[9] In this chapter, the types are discussed in a general context without applying them to a specific problem. An application of heterogeneous decision makers is presented in chapter 2 where real subjects participate in a dynamic choice experiment.

[10] In the literature there are different ways that these types have been modeled. In Al-Najjar and Weistein (2009) the resolute type coincides with the sophisticated as both are dynamic consistent and commit to their initial choice. Our definition of the sophisticated type is closest to Hammond and Zank (2014) who assume backward induction methodology that is not necessarily dynamic consistent as the decisions are made in a *myopic* way. This modelling way agrees with the sophisticated type in Hey and Panaccione (2011) with the difference being that no updating takes place. In addition, Hey and Panaccione (2011) distinguish between the naive and the myopic type (we use this term interchangeably for the same type) with the latter being the decision maker that is confused.

to reassure that there will be zero deviation from the *ex-ante* plan. Al-Najjar and Weistein (2009) characterise this strategy as *aversion to information* where the agent commits to not learning whether an event occurred or not.

1.4.2 Naive

A *naive* or *myopic* decision maker fails to understand the sequential nature of the problem. As a consequence, each of the stages is faced independently of the other, strategy that leads to dynamic inconsistencies and dominated results. The allocation at each stage is based on the optimisation of the objective function at the current stage, or stating in a different way, the decision maker solves a series of static problems and maximises present utility. When a decision maker faces a series of sequential problems, she fails to anticipate the impact that the current choices will have to the future utility and thus at each period maximises the current objective function. Naive decision makers ignore that they are time inconsistent since they tend to evaluate the several alternatives and to choose according to what seems to be optimal at present. As a result, the decisions that are made differ from those that had been planned.

1.4.3 Sophisticated

Strotz (1955-56) and later Pollak (1968), were among the first to recognise that the resolute type or in other words the pre-commitment strategy is not always the optimal decision. More specifically, the idea is that a decision maker that is not able to commit to her or his future behaviour, would prefer to adopt a strategy of consistent planning and then pick up the optimal plan that will be actually followed sketching the profile of a *sophisticated* type. A sophisticated decision maker uses *backward induction* in order to figure out the optimal strategy for every given problem. It is quite close as an idea to what Selten, 1965 has suggested where the problem can be defined as an extensive form game and each step is solved as a sub-game perfect equilibrium. The agent of this type uses *backward induction* in order to determine the optimal strategy that she is going to follow. Starting from the last decision nodes of the tree she is able to take into consideration all the alternative outcomes and their respective prospects. Applying the same approach to all the previous decision nodes she can define the optimal path that will lead her from the start of the tree to the most preferable node. Following this process she will be dynamically consistent and the *ex-ante* preferences will coincide with the *ex-post*. Hammond and Zank (2014) explain that this way of modelling has to do with the fact that the individual is not able to pre-commit. Using backward induction she concludes to an optimal plan of action for the whole problem that starts from period 1. Sophistication is the prevailing model in economics.

A similar categorisation among different types has been done before for risky choice by Machina (1989) who makes the differentiation among the different types of behaviour using the criterion of consequentialism. There are the so called *α-people*, the dynamically consistent agents that follow Expected Utility, the *β-people* that are non-Expected Utility agents and apply consequentialism acting in a time-inconsistent way (myopic behaviour), the *γ-people* who are non-Expected Utility agents but are dynamically consistent and finally the *δ-people* who are characterised as sophisticated. As is later shown, when the preferences collapse to Expected Utility, all types coincide as both consequentialism and dynamic consistency are satisfied.

1.5 Experimental Literature

1.5.1 Dynamic Decision Making Under Risk

As the experimental investigation of dynamic decision making under ambiguity is a relatively new field, we start by surveying the studies that include problems with exogenously given probabilities for the various states of the world (presence of risk). In this literature there are several experiments that have been conducted which focus either on testing the several hypotheses and axioms of the theoretical models of dynamic decision making under risk (mainly EU) or to discover the different strategies that the agents apply when they face inter-temporal problems in order to resolve dynamic inconsistency. We present the experimental findings on risky choice for two reasons. The first is that there is a lack of empirical studies on how do people update ambiguous beliefs or stating in a different way how do they decide in a similar dynamic problem. On the other hand, risk can be thought as an extreme form of ambiguity. In both cases, it is useful to present the various different experimental protocols that have been applied, their main theoretical

predictions and results and then to discuss how this analysis can be useful for the case of ambiguous beliefs. A famous result stated by Karni and Schmeidler (1991) states that non-EU subjects who violate the independence axiom, will necessarily violate one of the axioms of *dynamic consistency, consequentialism* or *reduction of compound lotteries*. A number of experiments aim to test this result.
Early experiments on dynamic decision making in the presence of risk, include the work of Tversky and Kahneman (1981). Cubitt et al. (1998) focus on testing several hypotheses of the Expected Utility theory, the violation of the axiom of independence and the consequences that this violation has to the other principles of dynamic choice (separability, timing independence, framing independence and reduction of compound lotteries). In the experimental design they use, they present to the subjects several different decision trees and test for the equivalence of the four dynamic choice principles. They also test for the common ratio effect. Their results report violation of the common ratio effect as well as rejection of the standard theoretical strategies used to explain this effect. These results are also verified by Hey and Paradiso (1999) and Hey and Paradiso (2006), though in a different framework. In Hey and Paradiso (1999), the issue of timing indifference is being examined. More specifically, they ask whether the timing of ambiguity resolution affects the decisions of the agent and whether a sequential choice problem differs from an equivalent planned one. They present an experiment with three equivalent choice problems, a sequential choice problem, a planned choice problem and a non-sequential choice problem, showing that the assumption of timing-independence is not appropriate for individual preferences. In Hey and Paradiso (2006), the authors collect data on the evaluations of the subjects for several decision problems that are strategically equivalent, but differ in respect of their temporal framing. They obtain the interesting finding that many of the subjects realise that they are dynamically inconsistent and are willing to pay an amount of money to avoid their inconsistencies being exposed.

A different approach on dynamic decision making under uncertainty has focused on the psychological factors that influence behaviour. Experimental evidence on this includes the work of Cubitt and Sugden (2001) and Busemeyer et al. (2000). In the first paper, the authors discuss the role that certain emotions play in the formation of decisions, and the importance of affective experiences in dynamic choice under risk. The main result is a rejection of the timing independence principle as well as rejection of the principles of separability. In Busemeyer et al. (2000) and later in a replication of the experiment in Johnson and Busemeyer (2001), the focus is on the psychological aspects of the problem and similar questions are applied in the presence of risk in order to test the way that emotional factors affect dynamic consistency. Their findings include violations of dynamic consistency but not of consequential consistency. The main drawback of the experiment is the lack of incentivised responses.

Additionally, there are the experimental studies that focus on *updating issues*. As the decisions concerned are under risk, the main test is constrained on the validity of the Bayesian updating rule and the various biases that may appear. Charness and Levin (2005), compare the Bayesian rule and the *reinforcement* heuristic. There are two states of the world and at each state either the white balls win or the black balls wins. In addition, there are two urns. One has balls of the same colour that are valuable or valueless balls, depending on the state and the other urn contains a mixture of the two. The subjects are getting paid for drawing valuable balls. After a draw they have the opportunity to change urn or to keep drawing from the same. They also controlled for emotions. They found that when the two heuristics are aligned, there is very low rate of switching error and people behave close to the Bayesian Expected Utility model. The same happens when affection is removed. When the two heuristics disagree, then there is a mixture of behaviour. Holt and Smith (2009), report an experiment where subjects are asked to assess probabilities for unknown events based on the information they gather from observing random draws of their risk representation device. As a device, they performed draws from two different cups that contained light and dark marbles in different proportions. Using the Becker-DeGroot-Marshak (Becker et al. (1964)) method to elicit beliefs, they test how well the Bayesian update rule can predict subjects' decisions. The experiment included two versions, a standard one in the lab and a web-based experiment. The results show that the reported probabilities appear to be slightly upward biased when the Bayesian prediction is low and similarly being downwards biased when the prediction is high, biases that become stronger for means than for medians. In addition, they find *representativeness bias* to be present.

Two important aspects that have been tested in this study include the way that different prior information affects the assessment of probabilities, and how the average reported probability changes when uniform signals are reinforced by subsequent draws. The former shows that when the priors are not equal (50% for each cup) then there is a tendency to predict too high for low probabilities (prior =1/3), behaviour that is not verified in case of high priors (prior =2/3). Regarding the representativeness bias, they found that the average elicited probabilities are higher than the Bayesian predictions when the draws are representative of one of the two cups and similarly elicited probabilities are driven down when the draw is in contrast to the expected. Charness et al. (2007) test both individuals and groups on whether they satisfy first order stochastic dominance and if they respect the Bayesian rule when decisions made under risk. The design is similar to Charness and Levin (2005) with the only difference being that they allow for interaction by forming small groups. Their main results include that when subjects make decisions in isolation, they tend to violate first order stochastic dominance, violation that is significantly reduced when the subjects are allowed to communicate and take common decisions in groups. They claim that part of the deviations from EU that are observed in experiments, are mostly due to the nature of the lab, where decisions are made in an unfamiliar, isolated environment. Poinas et al. (2012) report an experiment on updating under risk, where the subjects receive a message that constraints the possible states of the world. The subjects must determine the number of yellow balls, in an urn containing 20 balls that are either blue or yellow. They observe a sample of 6,10 or 14 balls. The subjects are asked to report the number of yellow balls, which also consists the prior beliefs. Then, they receive some signal on the composition of the urn, and subjects are again asked to predict the number of balls. They find that signals are helpful for more precise predictions. When signals confirm the initial belief, it is quite likely that subjects will change their prediction, when the level of risk is high. Recently, in a different type of study, Deryugina (2013) discusses how beliefs about climate change are formed and updated. Starting from the standard premise that both updating and learning are realised in a Bayesian way, she tests how well this model predicts behaviour and if there are biases (representativeness, availability and spreading activation) observed. Using a multi-year survey she tests how individuals form inferences about the occurrence of global warming. She found that some features of the updating process are Bayesian, but she also finds that the updating process is in line with the representativeness heuristic. There is also some evidence of the availability heuristic, where people give more weight to local temperatures. The majority of the studies presented above make the assumption that the decision makers are *risk neutral*. Antoniou et al. (2013) use a standard psychological experiment, suitably adapted to accommodate experimental economics practices, they account for preferences consistent with non-linear utility and they estimate structural models in order to test whether taking into consideration non-linearities, actually reduces deviations from the Bayesian updating or not. They applied two tasks, one to elicit risk preferences and the second to elicit beliefs. The elicitation of risk preferences was realised with the use of lottery pairs. To elicit beliefs, they used two boxes (blue and white), where in each urn there were a number of 10-side dice with coloured sides (blue and white). Then, subjects were given information on a roll of all the dice in the box, which allowed them to form a prior. After announcing the outcome of the rolling, subjects were asked to place their bets in 19 different betting houses. With this methodology, they were able to find the "switch point". They found modest risk aversion. They also found that failing to correct for non-linearity provides stronger support for the Bayesian rule.

1.5.2 Classifying Different Types

Then, another approach in the literature of dynamic decision making with exogenous probabilities, aims to test how individuals that are dynamic inconsistent resolve this inconsistency (when they do). The latter, was motivated by the seminal work of McClennen (1990) who firstly categorised three behavioural types: the resolute, the naive and the sophisticated decision maker. In Hey and Lotito (2009), the experimental structure is formed in such a way that allows the investigation of behaviour and the preferences of the agents at the same time. Using problems that were presented in both static and dynamic decision trees and with several combinations of outcomes and probabilities they were able detect the non-Expected Utility agents and furthermore to identify the specific type of behaviour that each subject adopts. An additional feature of the experiment was the use of a second-price sealed-bid auction in order to obtain the subjects' evaluations for the different trees. Several restrictions needed to be made regarding the preference functional, assuming that preferences can be represented by the *Rank Dependent Expected Utility* model with a CARA or CRRA utility function and a weighting function $w(p)$ following either the *Quiggin* or the *Power* specification. The results were interesting for verifying the dynamic consistency problem since 50% of the subjects were proven to be of naive type, 40% to be resolute and only 10% sophisticated. In another experiment, Hey and Panaccione (2011) applied a different experimental framework to test dynamic decision making and asked the subjects to allocate an amount of money m in N decision problems that included two-stage problems. The decision makers were assumed to have *Rank Dependent Expected Utility* preferences and a non-linear weighting function $w(p)$. This allowed for a classification of the subjects in different types (myopic, naive, sophisticated and resolute) with most of the subjects appearing to be resolute. More recently Nebout and Willinger (2014), aim a type-categorisation from data obtained from an experiment in dynamic choice under risk. They define three different types, naive, resolute and sophisticated and they distinguish between dynamic consistency and strategic dynamic consistency (the latter is the case where the choice at each decision node agrees with the strategy that has been chosen at the beginning of the problem). They use three different tasks and they proceed to two different tests. Initially, they divide subjects to EU and non-EU. Then, they test if they are strategically dynamic consistent or not. The results show that only 20% of the subjects are EU and 85% of the EU subjects satisfy DC and (92.5% satisfy strategic dynamic consistency-SDC). For non-EU subjects, they find that 72.5% satisfy DC and 65 % SDC. A serious drawback of their methodology is the low number of questions. They ask in total 16 questions which were divided in two groups: the parameter elicitation questions (11) and the categorisation questions (5). In addition, they do not explicitly explain how the lotteries were played out for real at the end of the experiment. In their analysis they have an interesting feature, where they do not exclude non-EU decision makers from having dynamically consistent preferences. Finally, Houser et al. (2004), in a different framework, present a methodologically interesting experiment, where instead of assuming different types and their respective optimal decisions, they classify the subjects based on a Bayesian type classification algorithm. They claim that this approach allows to draw inferences about both the type and the number of decision rules that are present in a population. The experiment is on a dynamic stochastic optimisation problem under certainty where subjects choose between two discrete alternatives in 15 time periods. Applying their algorithm to the data, identified three types, the near rational, the fatalist and the confused.

1.5.3 Dynamic Decision Making Under Ambiguity

The literature is extended to a lesser degree when dynamic problems under ambiguity are investigated. To the best of our knowledge, there are only two experiments that focus on pure[11] updating under ambiguity, Cohen et al. (2000) and Dominiak et al. (2012). Cohen et al. (2000) are motivated by the evidence on the non-universality of the Bayesian rule and the contradictory results on how ambiguous beliefs are updated. They focus on the descriptive validity of the main two updating rules, the *Maximum Likelihood* updating rule (MLU) and the *Full Bayesian* rule. Using the standard extension of the three-colour Ellsberg urn, with the revelation of some information at an interim point, they test for the two updating rules confirming the Ellsberg type behaviour and show that the FBU rule is used more often that the MLU rule. The drawback of the paper is that separability is assumed (an assumption close to consequentialism -the capacity v of the decision maker depends only on the available information and on possible counterfactuals). The latter does not allow for a direct test of which axioms subjects satisfy. In addition the experiment was not incentivised in monetary terms. Dominiak et al. (2012) using an experiment on the Ellsberg urn in a dynamic framework, try to test the two rationality criteria on dynamic decision making, dynamic consistency and consequentialism. The design they adopt is the same as in Cohen et al. (2000). The main result is that subjects adhere to C and reject DC (45%, contrary to 21% that satisfied DC.) . They also find support for the Full Bayesian updating rule. In order to avoid offering an option that expresses "indifference", they decided to ask subjects how confident they feel about their choices, using a measured scale (Null-Very Strong). They suggest that an answer close to null is equivalent to an expression of indifference regarding preferences. It is not quite clear why this

[11] The term pure is used, to distinguish this kind of experiments from those that use a sequential decision task, but they also involve some kind of learning. We expand on this later.

should be the case and in top of that, there was no incentive mechanism to induce subjects to reveal their true level of confidence (assuming that it can be measured). This experimental protocol suffers from a serious drawback. The results are based only on 4 observations per subject, regarding their preferences on the bets on the different urns. Even if one wants to allow for stochastics in the decision process, the number of questions is extremely low to draw any reasonable inferences from the data. The same applies to Cohen et al. (2000).

1.5.4 Learning under Ambiguity

A slightly different way of updating of prior beliefs under uncertainty, happens when the decision maker receives information repeatedly. If the case is such, then it is reasonable to expect that *learning* takes place. This learning process plays a significant role, leading to results closest to optimality. The intuition behind this lies on the fact that consecutive updating of beliefs will allow the agent to form a more precise subjective distribution function which will highly resemble the objective one, allowing the decision maker to get rid of any kind of ambiguity. The theoretical contributions on this subject include Epstein Moreno and Rosokha (2013) who report the results of an experiment where subjects were making sequential choices over pairs of lotteries involving two kinds of urns, a risky one and an ambiguous one. Black or white marbles were drawn with replacement from the urns and the subjects were asked to choose between a lottery and a certain amount of money. A series of draws was providing information on the composition of the urn, and the it was assumed that the evaluation of the bets was based on this information. Two models of learning are estimated, the Bayesian updating allowing for base rate fallacy and the reinforcement learning model. They find that the information incorporated in the ambiguous situation is lower compared to the risky one. Information is incorporated in a consistent way with Bayesian updating in the case of risk. In the experiment that Baillon et al. (2013) conducted, the authors aim to test the effect of learning on beliefs and ambiguity attitudes. They study the updating of decision weights in the context of a general preference model. As source of uncertainty, they used the variation of stock returns of IPOs (Initial Public Offerings) traded at the New York Stock Exchange. They included for three informational conditions, no information, one week information of the daily returns of the stock and one month information. Assuming a Power utility function and using nonlinear least squares, they estimate the parameters of interest. Although they claim they conduct an experiment on learning, in fact the experiment is just decision making with different levels of information. Each informational condition included different stocks. Consequently, there was no learning process where the subjects acquire new information, update their beliefs and then decide again. Finally, Qiu and Weitzel (2013) use an experimental framework that aims to elicit multiple priors. They used two different types of urns with type 1 containing 3 black and 6 white balls and type 2 containing 6 black and 3 white balls which were constructed based on numbers that the subjects have submitted in advance, without knowing where these numbers are going to be applied. Then, they asked subjects to estimate the probability that an urn of type 1 is chosen, with the use of a quadratic scoring rule[12]. This would form a prior. Afterwards, using again a scoring rule, the subjects were asked estimate the priors of the rest of the participants. This was applied in order to elicit the confidence regarding the distribution of multiple priors. In order to test the updating process, balls were drawn with replacement form the urn. This allows to update both the beliefs and the confidence towards the priors. The results do not show strong deviations from Bayesian learning. They also find that a confirmatory signal differs from a contradictory signal in terms of priors and distributions of multiple priors. Subjects under-react to confirmatory and over-react to contradictory signals.

1.6 Ellsberg Paradox with MEU Preferences

In this section, we use one of the most commonly used models of decision making under ambiguity and show an example of how non-dynamically consistent updating rules are applied in a sequential choice task. This example shows that depending on the preferences (and the updating rule that is applied), there is no unique way to update beliefs in ambiguous environments and consequently further empirical investigation is required. We consider the standard Ellsberg three-colour urn problem, with 1/3 of the balls to be Red and the rest 2/3 to be Blue and Yellow creating the state space $\Omega = \{R, B, Y\}$. Then, to introduce dynam-

[12] A weakness of this methodology is that risk neutrality is assumed.

ics, we assume a three-period model where at $t = 0$ the decision maker chooses between f' and g', at $t = 1$ receive partial information in the form of a filtration that presents the complementary event (e.g. $\mathcal{F}_1 = \{\{R, B\}, \{Y\}\}$ which provides the information whether the ball IS Yellow or not) and at $t = 2$ all relevant uncertainty is resolved. The MaxMin Expected Utility (MEU) model, proposed by Gilboa and Schmeidler (1989) assumes that an individual acts as if she had multiple (additive) priors regarding the subjective probability. The Expected Utility of a prospect is the minimum Expected Utility across all these priors. Then objective is to maximise across these minima. Given a set of priors $\mathcal{P}$, the utility of an act f over the set of priors $\mathcal{P}$ is given by:

$$V(f) = min_{p \in \mathcal{P}} E_p[f]$$

Using this notion and taking into consideration all the possible distributions of Blue and Yellow balls, then a set of priors is the following:

$v(R) = \frac{1}{3}, v(B) = v(Y) = 0$
$v(R \cup B) = v(R \cup Y) = \frac{1}{3}$; $v(B \cup Y) = \frac{2}{3}$
$v(S) = 1$

An alternative way to express the prior distribution is to write down all the possible distributions. For the three outcomes, and since $v(R) = \frac{1}{3}$, the set of priors is

$$\{P(R) = \frac{1}{3}, P(B) = p, P(Y) = \frac{2}{3} - p\} \tag{1.1}$$

with $p \in [0, \frac{2}{3}]$. This allows for a lower and an upper bound for p in this interval and can be written as $p_* \leq p \leq p^* \Rightarrow \frac{1}{3} - \lambda \leq p \leq \frac{1}{3} + \lambda$ with $\lambda \in [-\frac{1}{3}, \frac{1}{3}]$.

As lower bounds have been characterised, it is possible to represent the set of priors with the use of the Marschak-Machina Triangle. In Figure 1.1 the Marschak-Machina Triangle is illustrated, with the probabilities of Red and Blue to be depicted on the horizontal and vertical axis respectively. Assuming that λ takes values different than -1/3 or 1/3, which means some positive probability is attached to the other two states of the world, these priors lie along the AC line as the probability of Red is fixed to 1/3. Consequently, this gives two possible sets of priors $\{P(R), P(B), P(Y)\}$, one at point A ($\{1/3, \frac{1}{3} - \lambda, \frac{1}{3} + \lambda\}$) and one at point C ($\{1/3, \frac{1}{3} + \lambda, \frac{1}{3} - \lambda\}$). $\underline{P(B)}$ stands for the lower value that P(B) can take and is equal to $1/3 - \lambda$. Similarly, $\overline{P(B)}$ is the upper limit of $P(B)$ which is equal to $1/3 + \lambda$. Both the updating rules that we consider apply prior by prior updating to these sets of priors with the only difference that different updating rules focus on different priors. The crucial point now is to determine how these updated weights are formed. Gilboa and Schmeidler (1993) presented and axiomatised several update rules for both the cases where there is no unique additive prior and for non-additive probabilities.

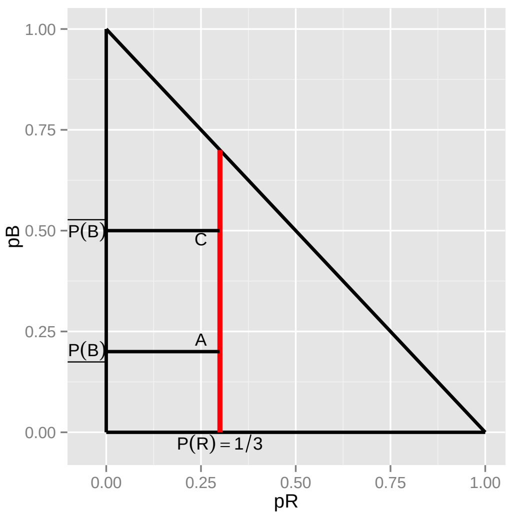

Figure 1.1 Priors in the Ellsberg Paradox

1.6.1 Updating Rules

The Maximum Likelihood Update rule (MLU) that has been axiomatised in Gilboa and Schmeidler (1993) and earlier by Dempster (1967) and Shafer (1976). Given that the ambiguous beliefs are characterised by a capacity v, given an event E, the beliefs that the state A has occurred is given by:

$$v_{\text{MLU}}(A|E) = \frac{v(A \cup E^c) - v(E^c)}{1 - v(E^c)} \tag{1.2}$$

for all $A \in \mathcal{A}$. According to this update rule, the decision maker considers only the prior that maximises the probability of the event and updates probabilities according to the Bayesian rule. This provides a new set of probabilities for the event that forms exactly the Core. This is a convex capacity, given that v is convex.

The Full Bayesian Update rule (FBU) is defined as:

$$v_{\text{FBU}}(A|E) = \frac{v(A \cap E)}{v(A \cap E) + 1 - v(A \cup E^c)} \tag{1.3}$$

for all $A \in \mathcal{A}$. In this rule, all the priors are updated in a Bayesian way and the set of posteriors is used to evaluate the different acts.

Let us now see how these rules apply in the case of the sequential Ellsberg urn. A ball is randomly drawn from the urn. Let a filtration be $\mathcal{F}_1 = \{R, B\}$ that is to say the information "the ball is not Yellow" is revealed to the decision maker. Now the objective is to choose between bets f and g. Using this partial information, the payoffs in Yellow are not relevant anymore and as a result only the priors of Red and Blue are updated. The initial set of priors at A and C, are now reduced to $\{1/3, \frac{1}{3} - \lambda\}$ and $\{1/3, \frac{1}{3} + \lambda\}$. First we consider the MLU rule. The event now is $R \cup B$ so the decision maker updates those priors that maximise the probability $P(R \cup B)$. In this example, this happens when the $v(B) = \frac{1}{3} + \lambda$ as $v(R)$ is constant and equal to $\frac{1}{3}$. In the triangle, the set of priors that satisfies this condition is the one in point C. Using 1.2 and expanding the capacities, we have:

$$v_{\text{MLU}}(R|R \cup B) = \frac{v(R \cup (R \cup B)^c) - v(R \cup B)^c)}{1 - v(R \cup B)^c)}$$

$$v(R \cup B)^c) = 1 - v(R \cup B) = 1 - v(R) - v(B) = \frac{1}{3} - \lambda$$

$$v(R \cup (R \cup B)^c) = v(R) + v(R \cup B)^c) = \frac{2}{3} - \lambda$$

$$1 - v(R \cup B)^c) = \frac{2}{3} + \lambda$$

and combining all the above:

$$v_{\text{MLU}}(R|R \cup B) = \frac{1}{2 + 3\lambda}$$

and similarly, the updated capacity for Red (always using the prior probability of Blue that maximises the event[13], the updated capacity for Blue is:

[13]In this example it is easy to see that the probability of the event is maximised when the prior of Blue is equal to $\frac{1}{3} + \lambda$, and since λ is a non-negative number it is clearly greater than the prior of Red ($\frac{1}{3}$). In our experiment, this is not clear as there are no fixed capacities and a check of which

$$v_{\text{MLU}}(B|R \cup B) = \frac{1+3\lambda}{2+3\lambda}$$

Applying the FBU rule, the set of priors that is updated depends on the colour that we focus on. As it was discussed before, and using Figure 1.1, we can identify two different sets of priors for $\{P(R), P(B)\}$ at points A and C which are respectively the sets $\{\frac{1}{3}, \frac{1}{3} - \lambda\}$ at point A and $\{\frac{1}{3}, \frac{1}{3} + \lambda\}$ at C. We can write the updated set as $P' = \{\frac{P(R)}{P(R)+P(B)}, \frac{P(B)}{P(R)+P(B)}\}$. This rule is applied to all the available sets of priors. Then the decision rule requires to choose according to the updated prior that provides the lowest Expected Utility. Take for instance a bet to a Red ball. The updated probability for Red is $\frac{1}{2-3\lambda}$ by updating the set of priors in A, or $\frac{1}{2+3\lambda}$ in C. Comparing the two, it is easy to observe that $\frac{1}{2+3\lambda}$ is the updated probability that minimises the Expected Utility. when the which offers minimises the Expected Utility, it is easy to see that. Updating both sets of priors The FBU rule suggests that this prior is updated using the Bayes rule, so writing this in the general form of the priors of expression 1.1, the set of posteriors will simply be , where the value of p depends on which vertex of the triangle we are in.

Assuming that we want to calculate the minimum Expected Utility for Red. Since A is not available anymore, this means that the set of priors is the set on C, $\{\frac{1}{3}, \frac{1}{3} + \lambda\}$. Substituting in 1.3, the posterior of Red conditional to the event $R \cup B$ when the set of priors that minimises the event (C) is updated is:

$$v_{\text{FBU}}(R|R \cup B) = \frac{1}{2+3\lambda}$$

as $v(R \cap (R \cup B)) = v(R) = \frac{1}{3}$ and $v(R \cup (R \cup B)^c) = v(R) + 1 - v(R \cup B) = \frac{2}{3} + \lambda$.

Similarly, the conditional probability for Blue that minimises Expected Utility, using the FBU rule, is calculated at the point B. This is the case as in B the set of priors is $\{\frac{1}{3}, \frac{1}{3} - \lambda\}$. In this case $v(R \cup (R \cup B)^c) = \frac{2}{3} - \lambda$ and substituting to the formula of the update rule the conditional probability is:

$$v_{\text{FBU}}(B|R \cup B) = \frac{1-3\lambda}{2-3\lambda}$$

Both update rules can be applied in the case of Hurwicz-type models (α-

prior is maximised must be done.

Expected Utility) using the formula:

$$v_i(A|E = \alpha w_i(A|E) + (1-\alpha)(1 - w_i(A^c|E))$$

with $i \in \{\text{MLU, FBU}\}$

1.7 Conclusion

This chapter surveyed the literature of dynamic decision making under ambiguity. After introducing some fundamental definitions of choice under ambiguity, we presented the main behavioural anomaly (Ellsberg paradox), due to which the ambiguity aversion literature emerged. Then, the way that the Ellsberg anomaly may appear in a sequential problem was illustrated, followed by a discussion on the extensions that the current models of static choice under ambiguity require, so that they can accommodate inter-temporal choice problems. Then, the relevant empirical evidence was discussed, based on the available experimental findings. Finally, using the example of the sequential version of the Ellsberg paradox, we briefly introduced updating rules for a special family of models and we show how different updating rules, lead to different results. The main message from this chapter is that, although there is a rich theoretical literature on dynamic choice under ambiguity that is continuously expanded, there is yet little evidence on what people are really doing when they face similar problems. Most of the theoretical contributions need to make extremely restrictive assumptions on the process that is followed when prior beliefs are updated. The latter, seems to have little connection with the actual cognitive procedure that is actually applied. In the next chapter, we present an economic experiment in which we allow for behaviour that deviates from the optimal. With the results gathered by the experiment, we aim to behaviourally test what people are doing when they face a sequential problem and they are not aware of an exact probability distribution regarding the various states of the world.

Chapter 2

Dynamic Decision Making under Ambiguity: An Experimental Approach[1]

2.1 Introduction

In this chapter we present the design and the results of an economic experiment on dynamic decision making under ambiguity. Building on the conclusions of chapter 1, it has become apparent that when subjects exhibit attitudes towards ambiguity (either aversion or preference for ambiguity) then preferences cannot be represented by the standard Expected Utility model. The direct consequence of this is that the decision will not necessarily be dynamically consistent, an assumption that is one of the main cornerstones in mainstream economic theory. In this experiment we aim to understand how people behave in a dynamic problem under ambiguity, how they update their prior beliefs and which is the alternative, non-EU model that best captures all these issues. One of the novelties of this experiment is that it deviated from the standard way that ambiguity is represented in the lab and instead of using standard Ellsberg type urns, a Bingo Blower was used, a device which eliminates the suspicion that the Ellsberg-type urns generate. Then we ask a series of allocation problems to the subjects in sequential choice

[1]This study was funded by the Research and Impact Support, Department of Economics, University of York (RIS 39), fund jointly awarded to John Hey and the author.

problem. Based on the data gathered from the experiment, we specify different types of decision makers and we fit the data to different preference functionals and updating rules. We identify three types of decision makers, the resolute, the naive and the sophisticated.

It is crucial to make a clarification regarding the term *dynamic*. While in the relevant literature the term dynamic has prevailed when decisions in similar environments are discussed, one should be very careful not to confuse it with the more generalised notion of dynamic choice which requires the decision to take place at different points of time (this may be the duration of one day, one week or even several years) and discounting of the future takes place. Alternatively, the term *sequential* choice seems to be more appropriate for the kind of problems that we discuss. It is a sequential decision task in the spirit that an agent is required to make a sequence of decisions under incomplete information based on the prior beliefs that somehow exist. After each decision, a certain amount of partial information is obtained, which allows for the updating of the prior beliefs. Based on the updated probabilities, a new decision must be made. The inconsistencies in the preferences that may be observed are due to the non-dynamically updating of the priors instead of some kind of time discounting.

To the best of our knowledge there are only two experimental studies, Cohen et al. (2000) and Dominiak et al. (2012), that address the issue of updating of prior beliefs in an sequential choice problem under ambiguity. Our study differs in four substantial ways from these studies. To begin with, the main differentiation is to be found in the representation of ambiguity. Both Cohen et al. (2000) and Dominiak et al. (2012), use the sequential Ellberg-type urn, as this was presented in chapter 1. Instead, in our framework, a *transparent* and *non-manipulable* device has been applied by using a Bingo Blower, in the same way as Hey and Pace (2014) did. A second difference is found in the decision task and the number of questions. While in the previous experiments, the decision task is constrained to pairwise choices, in our study we adopt allocation type questions, first introduced by Loomes (1991) and then exploited by Choi et al. (2007), Hey and Pace (2014) and Ahn et al. (2014), which seem to provide informationally richer datasets. Furthermore, we ask subjects a total of 60 questions compared to the just 4 questions that have been asked before. A third point has to do with the econometric specification. In order to proceed with our analysis, we make explicit assumptions on the utility function and the preference functional that characterise subjects' preferences. This allows us to investigate decisions by applying different theoretical models (and consequently test different models and updating rules). The final point that this study differs from what has already be done, is that it accounts for error in decision making. Assuming that subjects make their decisions with error, we account for this noise in the data by assuming a specific stochastic specification. The framework that we adopt allows us to consider non-linearity of the utility function.

As is highlighted by the results of Antoniou et al. (2013), failing to correct for the non-linearity of the utility function provides stronger behavioural support for the Bayesian rule. In this paper, they focus on decision making under ambiguity and they account for non-linearities in the utility function and the effects they have in deviations from Bayesian updating. In addition, they assume that the subjects are neutral towards ambiguity. They conclude, that in order to be able to capture ambiguity aversion, there are three extensions than need to be done, in theoretical, experimental and econometrics terms. The theoretical extensions include the use of preference functionals that allow for ambiguity aversion. Then, the extensions in the experimental framework and the subsequent econometric analysis require tasks that will make the identification of different theoretical structures possible. In this chapter, we apply all the above extensions. We aim to elicit beliefs by using an alternative method compared to the standard one that has been used in the literature. Using this kind of elicitation method, one is able to construct structural models that are capable of capturing attitudes towards ambiguity. The experimental framework, as is analytically described later, allows for a transparent and non-manipulable representation of ambiguous events. Finally, the stochastic specification along with the parametrisation of the preference functionals and combined with the fact that the analysis is done on a subject level analysis, allows for the identification and classification of several different structures that aim to model decision making.

The chapter is organised as follows. In the next section, the theoretical framework is discussed concerning theories of decision making under ambiguity and updating. Section 2.3 presents the experimental procedure. Then, in section 2.4

we present the econometric specification and the stochastic assumptions. In section 2.5 the results of the experiment are presented followed by an interpretation. Then we conclude. At the end of the chapter we provide two different appendices, one non-technical and one technical. The non-technical (Appendix C) includes the instructions for the experiment. Appendix A provides the analytical solutions for all the different types of decision makers and all the preference functionals that we test.

2.2 The Theoretical Framework

2.2.1 The Decision Task

An agent is endowed with an income m and is asked to make an allocation between three *Arrow* security assets ($x_s, s \in \{i,j,k\}$) where an Arrow security asset is defined as an asset that pays 1 monetary unit if the state of the world is s and zero otherwise. Defining as z a normalised payoff function of the capital allocated to asset x_s, the return on this asset is defined as:

$$z(x_s) = \begin{cases} 1 & \text{if } s \in E \\ 0 & \text{if } s \in E^c \end{cases}$$

There are three possible states of the world and the decision maker faces ambiguity about which state will occur. Although there is ambiguity, the agent is assumed to be always in a position to form some kind of *subjective* priors. At each state only one asset yields a payoff which is based on the exchange rate of this asset.[2]. We denote the three states of the world by the finite state space $\mathcal{S} = \{s_i, s_j, s_k\}$ where assets i, j and k imply payments in their respective states. The respective probabilities[3] for each state of the world are $p(s_i), p(s_j), p(s_k)$ which we simply write as p_i, p_j and p_k. Similarly, we define e_i, e_j, e_k as the exchange rates of the assets. The problem can be described with the use of a three-period model. At

[2]One may think of this exchange rate as the rate of return of this asset. In the experiment that we present later, this exchange rate was applied in order to transform experimental income to monetary value.

[3]Here we implicitly assume that these probabilities exist and are additive. We expand on this later.

$t = 0$, the agent makes an initial allocation to the three assets, based on her or his[4] subjective beliefs, the available income, the exchange rates and her individual characteristics (risk aversion, ambiguity aversion). The objective is to maximise her utility, knowing that only one state of the world will prevail and that she will get paid only by the asset that corresponds to this state. Before learning the actual state of the world, some partial information is revealed that allows the decision maker to *update* her initial beliefs and to change the initial choices if necessary. This means, that at $t = 1$ the information in the form of a filtration that presents the complementary event (e.g. $\mathcal{F}_1 = \{\{\text{i,j}\}, \{\text{k}\}\}$ is revealed, which provides the information of whether the state of the world is k or not. The filtration becomes trivial if the state is k. This additional information allows the decision maker to adjust the initial choices if necessary. This happens in the following way: assuming that the available information is that *the state of the world is not k*, which we will denote as $\neg k$, then the available income is reduced to $m - x_k$, where x_k is the amount that was allocated to asset k at time $t = 0$, and the decision maker is allowed to allocate this remaining income to the two remaining assets. At $t = 2$ all relevant uncertainty is resolved and the agent receives the payoff. The problem can be represented by a *decision tree* in the way that is defined by Raiffa (1968):

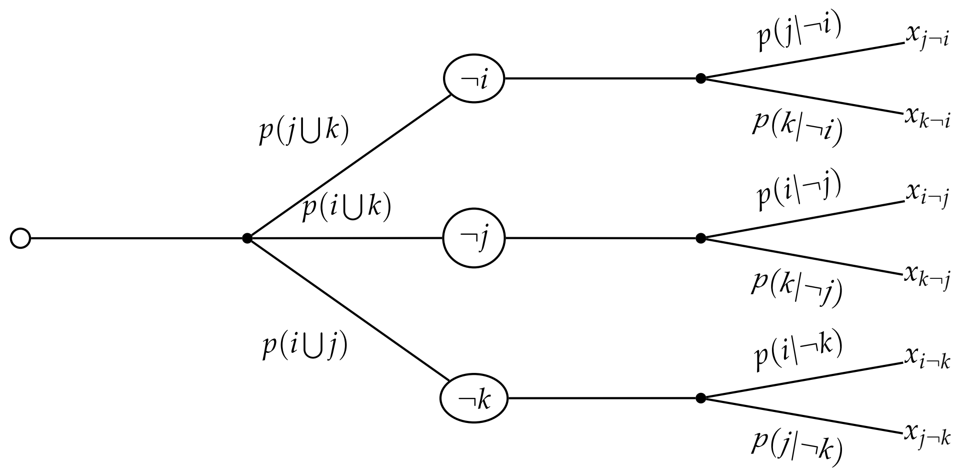

Figure 2.1 The 2-Stage Decision Tree

where $\neg i$ denotes that the state of the world *is not i*, $p(j \bigcup k)$ is the joint probability that the state is either j or k and is equal to the sum of the probability that the

[4] Henceforth, wherever in this thesis we want to refer to an agent, instead of "her or his" we will use "her" to include both, and avoid distracting repetition.

state is j and the probability that the state is k, $p(s_i) + p(s_k)$ (similarly are defined $p(i \bigcup k)$ and $p(i \bigcup j)$ and $p(j|\neg i)$ denotes the conditional probability of state j given that the state is not i. Based on the standard laws of conditional probabilities, this is equal to $p(j|\neg i) = \frac{p(s_j)}{p(s_j)+p(s_k)}$. Finally, $x_{j\neg i}$ is the conditional allocation to asset j when the state of the world is not i. This asset provides payoff equal to $e_j \times x_{j\neg i}$.

For the time being it is irrelevant what the source of ambiguity is. What is important is to define the decision criteria that the agent is using in order to make the allocation. We keep in line with the standard assumption that the objective of the decision maker is to maximise her utility levels based on some kind of preference functional. We now present the specifications that we will experimentally test. We include the Savage (1954) Subjective Expected Utility model (SEU), the Gilboa and Schmeidler (1989) MaxMin Expected Utility (MEU), the Schmeidler (1989) Choquet Expected Utility (CEU) and a parsimonious version of the CEU model, based on Kothiyal et al. (2014) which uses the Source method proposed by Abdellaoui et al. (2011). We call this model the Source Choquet Expected Utility model (SCEU).[5] . All the suggested preference functionals are accompanied with a specific story of how the ambiguous priors are updated with the arrival of new information. Based also on the fact that the problem can be represented as a decision-tree, it is reasonable to expect that a decision maker chooses paths along this tree when making a decision. This allows the specification of different types of behaviour towards those kinds of problems. Overall, combining the preference functionals, the updating rules and the types, provides a set of 14 specifications to be tested. In the next part we apply the following strategy. For each model, we present its main specification. Then, for each type[6] (Resolute, Naive and Sophisticated) we present the respective optimisation problem that the decision maker faces in its general form and we also present the updating process, if updating takes place.

[5]Although Machina (2013) shows that in the case of three outcomes, there are a series of paradoxes that the standard models are not able to accommodate. This is mainly due to the nature of the decision task which is constrained to pairwise choices. There is no straightforward way to see how this affects our case since we do not require strict ranking over the outcomes as the expression of indifference is allowed. We elaborate on this later.

[6]For an analytical presentation of how each types decides see chapter 1, section 1.4.

2.2.2 Subjective Expected Utility (SEU)

The SEU model is the standard model in the economics literature when one wants to model uncertainty. It assumes that the decision maker holds subjective beliefs on the various states of the world and these beliefs are represented by a unique, subjective, additive probability measure ($\sum p_s = 1,\ s \in \{i,j,k\}$). We assume that the agent derives utility from a standard Neumann-Morgenstern utility function[7] $u : X \to R$. Then, the optimisation problem of the decision maker is to find an allocation of the three different assets that maximises her utility and satisfies the budget constraint. Denoting the optimal allocation to assets i, j and k as x_i^*, x_j^*, x_k^* (which we will refer to as the *outcomes*), the problem in its general form can be written as:

$$\max_{x_i^*, x_j^*, x_k^*} \sum_{s=i}^{S} p_s u(e_s x_s) = p_i u(e_i x_i) + p_j u(e_j x_j) + p_k u(e_k x_k) \tag{2.1}$$

$$\text{s.t. } x_i^* + x_j^* + x_k^* = m \tag{2.2}$$

The Lagrangian writes:

$$\mathcal{L} = p_i u(e_i x_i) + p_j u(e_j x_j) + p_k u(e_k x_k) + \lambda(m - x_i - x_j - x_k)$$

where m is the available income and $e_s, s \in \{i,j,k\}$ the exchange rates. Based on the first order conditions and the budget constraint, it is possible to derive analytical formulas for the optimal allocations.[8]. In this model there is no way to capture *ambiguity aversion* as there is an implicit assumption that the subjective beliefs exist and sum to 1.

Different Types and Updating in SEU

An interesting property of SEU is that there is no way to distinguish between resolute, naive and sophisticated decision makers as all types update using the *Bayesian rule* which guarantees dynamic consistency. Nevertheless, in Appendix A we provide the solution for all the three types by assuming SEU preferences. A

[7]The form of the utility function is defined later.

[8]In Appendix A we derive the optimal allocations for all the specifications assuming a CRRA utility function. These specifications are then fitted to our experimental data.

resolute decision maker, solves the problem as if it is a one-period problem. The allocation that is made at stage 1 based on the subjective probabilities, coincides with the conditional allocations. The objective function is the same with Equation 2.1.

A *naive* decision maker fails to realise the sequential nature of the problem. At the first stage, this type behaves in the same way as a resolute does and solves the problem as if it is a static one. Then, at stage 2 receives the partial information, updates the prior beliefs according to Bayesian rule and then solves the maximisation problem that involves the two remaining states.

At stage 1 the decision maker makes the allocations based on the optimal at the present state (solves the problem as if it is a static one). Then receiving the partial information there are two effects. On the one hand, the part of the income that was allocated to the state that has not happened is lost. On the other hand, the initial beliefs on the different states are now updated based on this information. There are three events that can happen, that the state is not i, state is not j or the state is not k. We denote each *not* state as $\neg i, \neg j, \neg k$ for not i and so on. Let us focus on the case where the information received is that state of the world is not i. In the first stage, the problem is the same as in 2.1 which implies the vector of optimal allocations $x = (x_i^*, x_j^*, x_k^*)$. Knowing that the state of the world is not i, the remaining income m' is now equal to $m - x_i^*$. In addition, the prior beliefs for the two remaining states j and k are given by applying the Bayesian rule:

$$p(s_j|\neg i) = \frac{p(s_j)}{p(s_j) + p(s_k)}$$

$$p(s_k|\neg i) = \frac{p(s_k)}{p(s_j) + p(s_k)}$$

In the second stage, the naive decision maker behaves as if she is confronting a new problem that is not related to the previous one. The objective function in this case is

$$\max_{x^*_{j\neg i}, x^*_{k\neg i}} p_{j\neg i} u(e_j x_{j\neg i}) + p_{k\neg i} u(e_k x_{k\neg i}) \tag{2.3}$$

$$\text{s.t. } x^*_{j\neg i} + x^*_{k\neg i} = m - x^*_i \tag{2.4}$$

Where $x^*_{j\neg i}$ stands for the allocation to asset j when the state of the world is *not* i. Symmetrically are defined all the remaining 5 conditional allocations[9] as well as the solution to the objective function of the 2 remaining conditional states.

Finally, a *sophisticated* type solves the problem in two steps using *backward induction*. For the description of this type's strategy it will be useful to use the decision tree in Figure A.1. Solving backward, in the first step, the decision maker thinks what she would do if she was to reach a specific decision node. In this problem there are three decision nodes that can which form the set $\{s_{\neg i}, s_{\neg j}, s_{\neg k}\}$. At each decision node, the objective is to solve for the optimal conditional allocation taking as given the available income (which is the sum of the optimal amounts that were allocated at stage 1 of the experiment ($x^*_j + x^*_k$) if the state of the world is $s_{\neg i}$). This leads to the solution for the conditional allocations as a function of the available *conditional* income. We define as conditional income the available income at each conditional state and it is simply the initial income m minus the allocation to the foregone asset. If the state is not i, the conditional allocation $m_{\neg i}$ is equal to $m - x^*_i$. In the second step, taking into consideration these *conditional optimal allocations*, the decision maker solves for the unconditional optimal allocation x^*_i, x^*_j, x^*_k which is exactly what the decision maker does at the initial stage. In this task, the decision maker is liable to face three different events. This means that she has to solve three different conditional maximisation problem in the second stage, and the solutions from these problems will be used to define the initial allocations at stage 1.

Let us consider the case where the information revealed is *the ball is not i*. Then, the problem in the general form can be written as:

$$\max_{\{x_{j\neg i}, x_{k\neg i}\}} p(s_j|s_{\neg i})u(e_j x_{j\neg i}) + p(s_k|s_{\neg i})u(e_k x_{k\neg i})$$

$$\text{s.t. } x_{j\neg i} + x_{k\neg i} = x^*_j + x^*_k$$

which provides the solution for the optimal conditional allocations as a function of the optimal unconditional allocations (two additional problems need to be formed for the cases of not k and not j. We skip this part due to symmetry). In the second step, the decision maker solves for the optimal levels of the three

[9]In total 6 conditional allocations, 2 for each conditional state.

conditional incomes $m_{\neg i}, m_{\neg j}, m_{\neg k}$ taking as given the conditional allocations from step 1. In order to do so, one should solve the decision tree for the six conditional states, as viewed from $t = 0$. This requires the use of *compound probabilities* that provide the chances that specific outcome may happen. We recall that $p_{\neg i}$ stands for the unconditional probability that the state of the world is not i and is equal to $p_j + p_k$. With $q_{j\neg i}$ we denote the compound probability that state j will happen, when the state of the world is not i, which is simply the product of the unconditional probability that the state will be either j or k and the conditional probability of getting j when the event is that the state is not i. The formula for this compound probability is: $q_{j\neg i} = p_{\neg i} p_{j\neg i} \Rightarrow q_{j\neg i} = p_{\neg i} \frac{p_j}{p_{\neg i}}$. [10]. The objective function can now be written as:

$$\begin{aligned} \mathcal{L} = \frac{1}{2}[&q_{j\neg i}u(e_j x_{j\neg i}(m_{\neg i})) + q_{k\neg i}u(e_k x_{k\neg i}(m_{\neg i})) + q_{i\neg j}u(e_i x_{i\neg j}(m_{\neg j})) \\ &+q_{k\neg j}u(e_k x_{k\neg j}(m_{\neg j})) + q_{i\neg k}u(e_i x_{i\neg k}(m_{\neg k})) + q_{j\neg k}u(e_j x_{j\neg k}(m_{\neg k}))] \\ &+\lambda(2m - m_{\neg i}{}^* - m_{\neg j}{}^* - m_{\neg k}{}^*) \end{aligned} \tag{2.5}$$

The solution of the optimisation program in 2.5 provides the optimal levels of the conditional incomes. Manipulating these, it is possible to derive the optimal unconditional allocation that the agent applies when no information is available.

The full analytical solutions of how the types decide are provided in Appendix A. As was noted , the case of SEU preferences has no interest as the three types take exactly the same actions. The reason why we presented it in this specification, was to provide an illustration of how the problem is solved. This allows us to focus on the updating process. In the *non-expected utility* specifications that we will present in the rest of the section, the way that updating takes place, as well as the representation of beliefs, are totally different. Hence, it is possible to identify behavioural types of decision makers and accordingly to test for these differences.

[10]In this point becomes obvious how the different types coincide. Since the compound probability reduces to the unconditional one, it is reasonable to expect the same allocations, no matter the type of the decision maker.

2.2.3 MaxMin Expected Utility (MEU)

The MaxMin Expected Utility model was proposed by Gilboa (1987), Gilboa and Schmeidler (1989) and since then it has been extended and applied in several different frameworks (for an extended survey see Mukerji and Tallon (2004)). The MaxMin Expected Utility (MEU) model assumes that an individual acts as if she had multiple (additive) priors regarding the subjective probability. The expected utility of a prospect is the minimum expected utility across all these priors. Then objective is to maximise across these minima. Given a set of priors $\mathcal{P}$, the utility of an act f over the set of priors $\mathcal{P}$ is given by:

$$V(f) = min_{p \in \mathcal{P}} E_p[f]$$

In its general form, the model assumes that each decision maker constructs a set of probability measures on the state space S and this set constitutes her or his priors. The representation of the MEU model requires a $\succsim$ preference relationship on the set of actions. Then, when a series of standard Anscombe-Altmann axioms[11] along with two additional axioms, that of *C-Independence*[12] and that of *Uncertainty Aversion*[13] are satisfied, the representation theorem suggest that there exists a non-constant function $u : X \to \mathbb{R}$ and a non-empty, weak, compact, convex set $C \subseteq \Delta(\Sigma)$ of probability measures such that, for all $f, g \in \mathcal{F}$:

$$f \succsim g \Leftrightarrow \min_{P \in C} \int_S \left(\sum_{x \in \sup f(s)} u(x) f(s) \right) dP(s) \geq \min_{P \in C} \int_S \left(\sum_{x \in \sup g(s)} u(x) g(s) \right) dP(s)$$

Both in the decision problem that we present now and later when analysing the data from the experiment, a tractable version of this model is needed to model the behaviour of the subjects that comply with MEU preferences. However, the theory does not provide any kind of information on how this set of priors is formed. Since there are three states of the world, we follow Hey et al. (2010) and characterise the set of prior using the *Marschak-Machina Triangle (MMT)*, which consists a possible

[11] Weak order, continuity, monotonicity, non-triviality.

[12] C-Independence requires that for every action $f, g \in \mathcal{F}$ every constant $h \in \mathcal{F}$ and every $\alpha \in (0,1)$, $f \succsim g$ if and only if $\alpha f + (1-\alpha)g \succsim \alpha g + (1-\alpha)h$.

[13] For every $f, g \in \mathcal{F}$, if $f \sim g$ then for every $\alpha \in (0,1)$, $\alpha f + (1-\alpha)g \succsim f$.

way to represent priors.

We assume that there is a set of priors P and the three states of the world i, j, k such that:

$$\Pi = \left\{ P : P(\mathrm{i}) \geq \underline{p_i}, P(\mathrm{j}) \geq \underline{p_j}, P(\mathrm{k}) \geq \underline{p_k} \right\}$$

where $\underline{p_s}, s \in i, j, k$ is the lower bound probability for event s with $\underline{p_s} \geq 0$ and $\sum_s^S \underline{p_s} \leq 1$[14]. Using these priors that belong to the three-parameter space it is now possible to model the decision process of the subjects using the MEU model. It is possible to characterise these priors in the probability simplex using the Marschak-Machina Triangle. The triangle is a right-isosceles triangle with sides equal to 1. On the horizontal axis the probability for state i is represented while on the vertical axis for state k. The distance from the hypotenuse stands for the residual probability of state j. Figure 2.2 illustrates this idea. The small triangle inside the MMT is formed by the lower bounds of the priors for each of the states of the world.

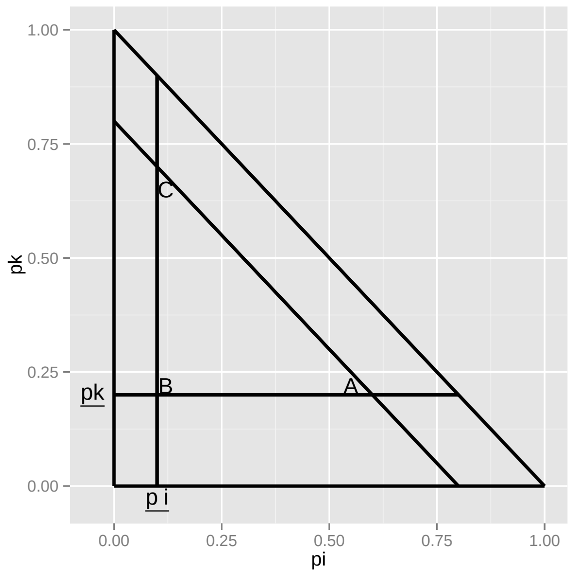

Figure 2.2 Prior Beliefs

The objective of the decision maker is, given the set of priors, to *maximise the minimum level of utility* that can be achieved. As the priors now are depicted as a triangle area inside the MMT (if the area of this triangle is small enough, such that it is transformed to a point, the MEU reduces to SEU), the utility is minimised along one of the corners of the triangle. Consequently, when calculating the MEU,

[14]When $\sum_s^S \underline{p_s} = 1$ the model is reduced to the SEU.

it is important to rank the different outcomes from the worst to the best. This means that when a decision is made, it is evaluated at its worst possible scenario. Consequently, depending on the various rankings, the utility is minimised at one of the three corners of the triangle (denoted A,B and C). Notice that each one of the prior sets, consists of an additive probability distribution. Table 2.1 shows the prior sets that characterise each of the vertices.

Table 2.1 Sets of Prior Beliefs

Case	**Weights**		
	w_i	w_j	w_k
A	$1-\underline{p_j}-\underline{p_k}$	$\underline{p_j}$	$\underline{p_k}$
B	$\underline{p_i}$	$1-\underline{p_i}-\underline{p_k}$	$\underline{p_k}$
C	$\underline{p_i}$	$\underline{p_j}$	$1-\underline{p_i}-\underline{p_j}$

Now returning to the sequential decision problem, the objective for a *resolute* decision maker is, based on the prior sets in Table 2.1 and the person-specific preference parameters, to find the vector of optimal allocations$x = \{x_i^*, x_j^*, x_k^*\}$ that maximises the lowest expected utility possible. Or stating it mathematically, by characterising the three different priors as $P_A, P_B, P_C \in \mathcal{P}$, to find the optimal vector x that solves 2.2.3[15].

$$\max_{\{x_i^*, x_j^*, x_k^*\}} (\min(EU(P_A, x), EU(P_B, x), EU(P_C, x))) \tag{2.6}$$

The interesting cases are to be found on the two other types (naive and sophisticated) where a process of information and revision of beliefs takes place. Although this family model consists of a generalisation of the Bayesian-type models, in contrast to the SEU, the Bayesian rule cannot be applied in a direct way and consequently dynamic consistency is not guaranteed. In the case of multiple priors, there are many different updating rules that have been suggested (see Dempster (1967), Dempster (1968), Shafer (1976) and more recently Jaffray (2008)) for situations where there are non-additive probabilities and a lack of a unique additive prior. The two that we consider here are based on the work of Gilboa and Schmeidler (1993) who axiomatise these Bayesian-type generalisation rules[16],

[15]The full expression of the objective function is provided in Appendix A.

[16]Gilboa and Schmeidler (1993) refer to these rules as *pseudo-Bayesian* rules.

the *Full Bayesian* updating rule and the *Maximum Likelihood* updating rule. In all these rules, it is explicitly assumed that the axiom of *consequentialism* is satisfied (thus dynamic consistency is violated if the preferences are not SEU, Ghirardato (2002)). The two rules are defines as following: given that the ambiguous beliefs are characterised by a capacity[17] v, given an event E, the beliefs that the state A has occurred, the Maximum Likelihood updating (MLU) rule is given by:

$$v(A|E) = \frac{v(A \cup E^c) - v(E^c)}{1 - v(E^c)}$$

for all $A \in \mathcal{A}$.

The Full Bayesian updating rule (FBU) is defined as:

$$v(A|E) = \frac{v(A \cap E)}{v(A \cap E) + 1 - v(A \cup E^c)}$$

where E^c stands for the complement of the event E. In chapter 1, we presented the way that these rules are applied to a particular case of the Ellsberg urn. Although it is a similar 3-state problem, it is fundamentally different in the fact that in the Ellsberg case, the probability of one state of the world is known, leaving only two uncertain states. In our decision task, there is ambiguity in all three states. Later we will show how this affects our result. Let us now see how beliefs are updated in the MEU framework.

As was described before, when some partial information is revealed, different priors are updated. For example, when the underlying event is that the state of the world is not i the decision maker updates her beliefs on j and k. Geometrically, this means that upon the reception of information, the dimension of the space that represents beliefs, is reduced from two to one and the beliefs are represented by a line instead of a triangle. This is illustrated in Figure 2.3.

Using the triangle, it is easy to illustrate the points that are updated. For instance, if the information revealed is that the state is not j then automatically the point B is not any more available as this is the point where the event j is maximised. In other words, the available points now are those along the line AC

[17]The notion of a capacity is fully described in the next section where the Choquet Expected Utility is presented. In general a capacity is a set function to [0,1] that represents beliefs but does not require to be additive.

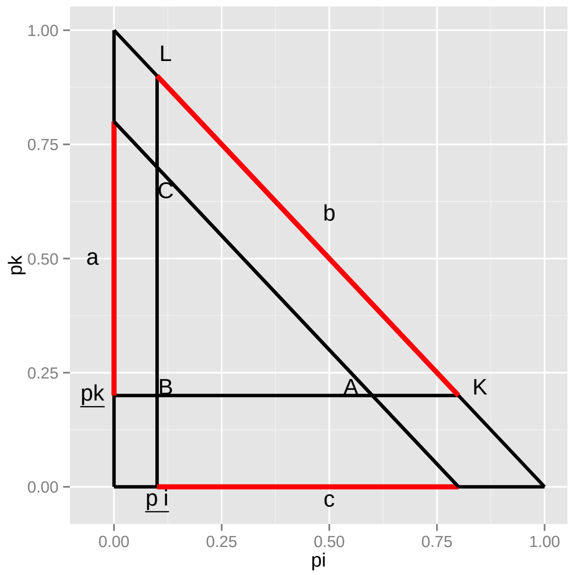

Figure 2.3 Feasible Update Intervals

(and the priors between A and C). Then again, using one of the two rules, each of the priors is updated (both the prior at A and at C). Then, the allocation between i and k is calculated by maximising the minimum expected utility inside the interval of the updated posteriors (this point lies along the segment AC and must be one of the two extremes).

Take for example the case that the event is *not* j. This means that from the triangle, we move to the side where $p_j = 0$ and this is true along the hypotenuse (side b). Consequently, the new updated sets will lie along b. However, notice that the updated beliefs cannot include the total length of this side as there are constraints imposed by the lower bounds of the beliefs. In our example, the feasible interval is the line between K and L. The interval cannot lie below K, as this corresponds to the level for the minimum prior for state k. Similarly, it cannot exceed L as this would require an updated probability for state i lower than its minimum value. The feasible interval is indicated with a red, bold line in Figure 2.3. In the same way, the feasible intervals for the other two conditional states are illustrated.

We return to out sequential decision task. Let for example the filtration $\mathcal{F}_1 = \{\{i,j\}\}$, that is the ball is not k. Then the MLU rule, requires that only the priors that maximise the probability of the event should be updated. In this case the priors that maximise the joint probability $P(i \cup j)$. It is to show that this happens at points A and B[18]. Consequently, instead of having only one unique posterior (as in the case of the Ellsberg urn) in this case the updating process provides an interval of possible posteriors. Practically, this means that the two rules coincide and thus, there is no way to discriminate between the two as they will result to the same decision. In section 2.2.3, we show that this holds, as in this framework the decision variable is continuous and as such there is only a unique allocation, given the interval of posteriors, that minimises the Expected Utility. In Tables 2.2 and 2.3, we summarise the set of posteriors for each possible event for MLU and FBU respectively.

Table 2.2 Maximum Likelihood Updated Posteriors

Event	Maximise	Points	Updated Posterior
not i	$P(j \cup k)$	B, C	$P'_B = [\frac{1-\underline{p}_i-\underline{p}_k}{1-\underline{p}_i}, \frac{\underline{p}_k}{1-\underline{p}_i}], P'_C = [\frac{\underline{p}_j}{1-\underline{p}_i}, \frac{1-\underline{p}_i-\underline{p}_j}{1-\underline{p}_i}]$
not j	$P(i \cup k)$	A, C	$P'_A = [\frac{1-\underline{p}_j-\underline{p}_k}{1-\underline{p}_j}, \frac{\underline{p}_k}{1-\underline{p}_j}], P'_C = [\frac{\underline{p}_i}{1-\underline{p}_j}, \frac{1-\underline{p}_i-\underline{p}_j}{1-\underline{p}_j}]$
not k	$P(i \cup j)$	A, B	$P'_A = [\frac{1-\underline{p}_j-\underline{p}_k}{1-\underline{p}_k}, \frac{\underline{p}_j}{1-\underline{p}_k}], P'_B = [\frac{\underline{p}_i}{1-\underline{p}_k}, \frac{1-\underline{p}_i-\underline{p}_k}{1-\underline{p}_k}]$

Table 2.3 Generalised Bayesian Update Rule

Event	Updated States	Updated Posteriors
not i	j,k	$P'_A = [\frac{\underline{p}_j}{\underline{p}_j+\underline{p}_k}, \frac{\underline{p}_k}{\underline{p}_j+\underline{p}_k}], P'_B = [\frac{1-\underline{p}_i-\underline{p}_k}{1-\underline{p}_i}, \frac{\underline{p}_k}{1-\underline{p}_i}], P'_C = [\frac{\underline{p}_j}{1-\underline{p}_i}, \frac{1-\underline{p}_i-\underline{p}_j}{1-\underline{p}_i}]$
not j	i,k	$P'_A = [\frac{1-\underline{p}_j-\underline{p}_k}{1-\underline{p}_j}, \frac{\underline{p}_k}{1-\underline{p}_j}], P'_B = [\frac{\underline{p}_i}{\underline{p}_i+\underline{p}_k}, \frac{\underline{p}_k}{\underline{p}_i+\underline{p}_k}], P'_C = [\frac{\underline{p}_i}{1-\underline{p}_j}, \frac{1-\underline{p}_i-\underline{p}_j}{1-\underline{p}_j}]$
not k	i,j	$P'_A = [\frac{1-\underline{p}_j-\underline{p}_k}{1-\underline{p}_k}, \frac{\underline{p}_j}{1-\underline{p}_k}], P'_B = [\frac{\underline{p}_i}{1-\underline{p}_k}, \frac{1-\underline{p}_i-\underline{p}_k}{1-\underline{p}_k}], P'_C = [\frac{\underline{p}_i}{\underline{p}_i+\underline{p}_j}, \frac{\underline{p}_j}{\underline{p}_i+\underline{p}_j}]$

We now describe the way that a *sophisticated* decision maker chooses. Starting from period t=1 and considering the case where the available information is that the state is not j. The allocation is now made between i and k. The decision maker should update the beliefs according to the available information and using this beliefs, to maximise the minimum expected utility. As was described before, the posterior beliefs now represented by an interval, instead of a single point, where the two extremes are defined by the updated priors $\frac{1-\underline{p}_i-\underline{p}_k}{1-\underline{p}_j}$ and $\frac{\underline{p}_i}{1-\underline{p}_j}$. The objective

[18] By definition, the MLU rule is applied to the priors that maximise the probability of the event. In this example this is the probability that the event is either i or j so we need to find at which sets of priors the $P(i \cup j)$ is maximised. At point A ,this probability is equal to $\underline{p}_i + \underline{p}_j$ while at points B and C it is equal to $1 - \underline{p}_k$ where the probability of the event is maximised. To see why this is true, assume that the event is maximised at point A. This means that $\underline{p}_j + \underline{p}_j > 1 - \underline{p}_k$ .This holds if and only if $\underline{p}_i + \underline{p}_j + \underline{p}_k > 1$ which cannot hold. In the case where this holds with equality, the model collapses to the Subjective Expected Utility and the little triangle collapses to a point.

function can now be written as:

$$\max_{\{x_{i\neg j}, x_{k\neg j}\}} \{\min\{w_i(u(x_{i\neg j}) + (1-w_i)u(x_{k\neg j}), w_k(u(x_{i\neg j}) + (1-w_k)u(x_{k\neg j})\}\}$$
$$\text{s.t. } x_{i\neg j} + x_{k\neg j} = x_i^* + x_k^* \tag{2.7}$$

where $w_i = \frac{1-p_i-p_k}{1-\underline{p_j}}$, $w_k = \frac{p_i}{1-\underline{p_j}}$, $x_{i\neg j}$ is the allocation to asset i at t=1 when the state is not j and x_i^* is the allocation to asset i at period t=0. Solving this, will provide an optimal conditional allocation which is a function of the initial allocations as these define the available income at the next stage. So we have $x_{i\neg j} = x_{i\neg j}(x_i^* + x_k^*)$ or writing the available income as $m - x_j^*$ the optimal allocation can be expressed as $x_{i\neg j} \equiv x_{i\neg j}(m - x_j^*)$. Consequently, $x_{k\neg j} \equiv m - x_j^* - x_{i\neg j})$.

$$\max_{\{x_{i\neg j}, x_{k\neg j}\}} \{\min\{w_i(u(x_{i\neg j}) + (1-w_i)u(x_{k\neg j}), w_k(u(x_{i\neg j}) + (1-w_k)u(x_{k\neg j})\}\}$$
$$\text{s.t. } x_{i\neg j} + x_{k\neg j} = m - x_j^* \tag{2.8}$$

The same process is applied for the two remaining conditional states. When the state is not i, the allocation is between assets j and k and the available income is $m - x_i^*$. This gives the allocation $x_{j\neg i} = x_{j\neg i}(m - x_i^*)$ and $x_{k\neg i} = m - x_i^* - x_{j\neg i})$. And similarly for the event not k, $x_{i\neg k} = x_{i\neg k}(m - x_k)$ and $x_{j\neg k} = m - x_k^* - x_{i\neg k}$ or the latter can be written as $x_{j\neg k} = x_i^* + x_j^* - x_{i\neg k})$. All the conditional allocations are a function of the optimal unconditional allocations. In Table 2.4, the decisions at $t = 1$ are summarised[19].

Table 2.4 Conditional Allocations at $t = 1$

State	**Conditional Allocation**
not i	$x_{j\neg i}(x_i^*)$ $x_{k\neg i}(x_i^*)$
not j	$x_{i\neg j}(x_j^*)$ $x_{k\neg j}(x_j^*)$
not k	$x_{i\neg k}(x_i^* + x_j^*)$ $x_{j\neg k}(x_i^* + x_j^*)$

At the second stage of the solution, the decision maker takes into consideration

[19]Notice that the optimal unconditional allocation for asset k is calculated as the residual $m - x_i^* - x_j^*$.

these optimal conditional allocations and solves for the optimal unconditional allocation at the stage where she holds no information. That is to say, at the second level, the decision maker solves the problem for x_i^* and x_j^*. The objective now is to solve the problem as it is conceived at $t = 0$ for all the six conditional allocations . The objective function is:

$$\max_{x_{j\neg i}, x_{k\neg i}, x_{i\neg j}, x_{k\neg j}, x_{i\neg k}, x_{j\neg k}} \{\min\{\text{P(not } i) \times \text{EU(not } i) + \text{P(not } j) \times \text{EU(not } j) + \text{P(not } k) \times \text{EU(not } k)\}\}$$

where $P(\neg i)$ stands for the joint probability that either j or k will happen. As the beliefs are ambiguous and we are not aware where exactly in the triangle these beliefs lie, we follow the same methodology as in the case of the resolute MaxMin preferences. For each point in the triangle, the MaxMin expected utility is calculated given the set of priors and the respective updated posteriors. The joint probabilities that an event will happen is summarised in Table 2.5 for each of the prior sets.

Table 2.5 Joint Probabilities for each Prior

Point	**not** i	**not** j	**not** k
A	$p_j + p_k$	$1 - p_j$	$1 - p_k$
B	$1 - p_i$	$p_i + p_k$	$1 - p_k$
C	$1 - p_i$	$1 - p_j$	$p_i + p_j$

We calculate the expected utility at all three points. For example, at point A, the optimisation problem can be written as:

$$\max_{x_{j\neg i}, x_{k\neg i}, x_{i\neg j}, x_{k\neg j}, x_{i\neg k}, x_{j\neg k}} \{\min\{P_A(\text{not } i) \times \text{EU(not } i) + P_A(\text{not } j) \times \text{EU(not } j) + P_A(\text{not } k) \times \text{EU(not } k)\}\}$$

and substituting:

$$\begin{aligned} EU_A =& (\underline{p_j} + \underline{p_k}) \times [(wu(x_{j\neg i}(x_i^*)) + (1-w)u(x_{k\neg i}(x_i^*))] + (1 - \underline{p_j}) \times [wu(x_{i\neg j}(x_j^*)) + \\ & (1-w)u(x_{k\neg j}(x_j^*))] + (1 - \underline{p_k}) \times [wu(x_{i\neg k}(x_i^* + x_j^*)) + (1-w)u(x_{j\neg k}(x_i^* + x_j^*))] \end{aligned} \tag{2.9}$$

where w is the posterior probability as was defined in Table 2.2 and differs at each conditional state. Then following the same line of thinking, the expected utility at points B and C are respectively:

$$\begin{aligned} EU_B =&(1-\underline{p_i}) \times [(wu(x_{j\neg i}(x_i^*)) + (1-w)u(x_{k\neg i}(x_i^*))] + (\underline{p_i}+\underline{p_k}) \times [wu(x_{i\neg j}(x_j^*))+ \\ &(1-w)u(x_{k\neg j}(x_j^*))] + (1-\underline{p_k}) \times [wu(x_{i\neg k}(x_i^*+x_j^*)) + (1-w)u(x_{j\neg k}(x_i^*+x_j^*))] \end{aligned} \tag{2.10}$$

$$\begin{aligned} EU_C =&(1-\underline{p_i}) \times [(wu(x_{j\neg i}(x_i^*)) + (1-w)u(x_{k\neg i}(x_i^*))] + (1-\underline{p_j}) \times [wu(x_{i\neg j}(x_j^*))+ \\ &(1-w)u(x_{k\neg j}(x_j^*))] + (\underline{p_i}+\underline{p_j}) \times [wu(x_{i\neg k}(x_i^*+x_j^*)) + (1-w)u(x_{j\neg k}(x_i^*+x_j^*))] \end{aligned} \tag{2.11}$$

Finally, the objective is to solve for the optimal levels of (x_i^*, x_j^*) that solve the following expression:

$$\max\{\min\{EU_A(x_i^*, x_j^*), EU_B(x_i^*, x_j^*), EU_C(x_i^*, x_j^*)\}\}$$

providing the optimal conditional and unconditional allocations.[20]

A *naive* decision maker combines the strategies of the resolute and the sophisticated type but in a reverse order. Firstly decides as if the problem is a one stage task by solving Equation . Then she receives partial information, updates her beliefs and uses these revised beliefs to solve Equation 2.7. Based on the type[21] of information revealed, the decisions may satisfy dynamic consistency or they may not.

FBU Vs MLU

In cases where the decision is constrained to discrete choice problems (such as pairwise problems), the distinction of the two rules is important as using either

[20]The optimal allocations for the MEU model were obtained numerically. For this purpose, the R package *optim* was applied using a general-purpose optimisation routine based on Nelder Mead, quasi-Newton and conjugate-gradient algorithms. This routine allows box-constrained optimisation and uses simulated annealing.

[21]The type of information can be classified as *good* news, *bad* news and *neutral*. This is a function of the probability that an event will happen and the respective exchange rate of this event. If for example there is a state that has high prior probability and yields a high payoff, by learning that this state has not happened, it can be regarded as bad news.

may lead to different choices (see sequential Ellsberg urn in chapter 1). If this is the case, then it is possible to identify and then test which of the rules is applied. In the problem that we are studying, due to the lack of a fixed probability this does not hold any more and consequently we are unable to identify which rule governs. We prove this formally. Consider again Figure 2.3. In this MMT p_i is depicted along the horizontal axis and p_k along the vertical and the internal triangle is formed by the lower bounds of $\underline{p}_i, \underline{p}_j, \underline{p}_k$. We assume that $0 < \underline{p}_i + \underline{p}_j + \underline{p}_k < 1$ [22] and also we assume that the information that the state is not k is revealed. The priors at each of the vertices A,B and C are:

$$\begin{array}{ll} \text{A} & [1-\underline{p}_j-\underline{p}_k, \underline{p}_j, \underline{p}_k] \\ \text{B} & [\underline{p}_i, 1-\underline{p}_i-\underline{p}_k, \underline{p}_k] \\ \text{C} & [\underline{p}_i, \underline{p}_k, 1-\underline{p}_i\underline{p}_j] \end{array}$$

Updating prior by prior using the Bayesian rule provides the posteriors A′,B′ and C′

$$\begin{array}{ll} \text{A}' & [\frac{1-\underline{p}_j-\underline{p}_k}{1-\underline{p}_k}, \frac{\underline{p}_j}{1-\underline{p}_k}] \\ \text{B}' & [\frac{\underline{p}_i}{1-\underline{p}_k}, \frac{1-\underline{p}_i-\underline{p}_k}{1-\underline{p}_k}] \\ \text{C}' & [\frac{\underline{p}_i}{\underline{p}_i+\underline{p}_j}, \frac{\underline{p}_j}{\underline{p}_i+\underline{p}_j}] \end{array}$$

The question now is whether the point C′ (the one that is excluded by the MLU rule) is within the set B′ to A′. Note that A′ is to the right of B′. We have that C′ is to the left of B′ if:

$$\frac{\underline{p}_i}{\underline{p}_i+\underline{p}_j} < \frac{\underline{p}_i}{1-\underline{p}_k}$$

that is if $1-\underline{p}_i < \underline{p}_i + \underline{p}_j$ which means that $1 < \underline{p}_i + \underline{p}_j + \underline{p}_k$ which is impossible. Also C′ is to the right of A′ if

$$\frac{\underline{p}_i}{\underline{p}_i+\underline{p}_j} > \frac{1-\underline{p}_j-\underline{p}_k}{1-\underline{p}_k}$$

that is if

$$\underline{p}_i(1-\underline{p}_k) > (1-\underline{p}_j-\underline{p}_k)(\underline{p}_i+\underline{p}_j) \Rightarrow \underline{p}_i - \underline{p}_i\underline{p}_k > \underline{p}_i + \underline{p}_j - \underline{p}_i\underline{p}_j - \underline{p}_j^2 - \underline{p}_i\underline{p}_k - \underline{p}_j\underline{p}_k$$

[22]We assume strict inequality as, if the sum of the probabilities is equal to zero then the problem is trivial and if it is equal to 1 it reduces to the SEU case.

which holds when

$$\underline{p_j} - \underline{p_i}\underline{p_j} - \underline{p_j}^2 - \underline{p_j}\underline{p_k} < 0$$

and simplifying

$$1 - \underline{p_i} - \underline{p_j} - \underline{p_k} < 0$$

which leads to contradiction. Thus, the sets should be the same. ■

The previous result indicates that the two updating rules will lead to the same allocations that the MEU preferences generate and even if the prior that the GBU rule requires to be updates is ignored, we will obtain the same result.

2.2.4 Choquet Expected Utility (CEU)

The Choquet Expected Utility (CEU) model belongs to the family of *rank-dependent utility* models and was firstly axiomatized by Schmeidler (1986), Schmeidler (1986) and later by Gilboa (1987) and Wakker (1989) among others. The main characteristic of this model is that beliefs do not need to be additive any more, as now these beliefs are represented by a *non-additive* capacity. Following Gilboa (2009), an unknown coin can be an example of a non-additive capacity. When the coin is *fair*, then it is quite reasonable to assume that $v(Heads) = v(Tails) = .5$. When one is not sure, the theory of non-additive capacities suggests that one can hold beliefs that do not sum up to 1. In other words, there is no contradiction in believing that $v(Heads) = v(Tails) = .4$ while at the same time $v(Heads \cup Tails) = 1$.

This means that the standard result $v(A \cup B) + v(A \cap B) = v(A) + v(B)$, when A and B are disjoint, does not hold any more and as we see later, this lack of equality is used to represent *attitudes towards ambiguity*. This non-additive capacity v is defined as a real-valued, non-linear increasing function from $S \to [0,1]$. This function satisfies the following properties:

1. $v(0) = 0$
2. $A \subset B$ implies $v(A) \leq v(B)$
3. v(S)=1

When 1-3 are satisfied at the same time, the function v consists a non-additive probability, or a capacity. One of the contributions of this model is that it allows the expression of *ambiguity aversion*. More specifically, Schmeidler shows that a CEU decision maker is ambiguity averse[23], when the capacity that expresses her beliefs is *convex*. A convex capacity for all events $A, B \in \Sigma$ satisfies:

$$v(A \cup B) + v(A \cap B) \geq v(A) + v(B)$$

Schmeidler's idea was based on the notion of an integral for capacities in the same why that it was developed by Choquet. More specifically, for a given capacity space (S, Σ, v) and a Σ-measurable function $\alpha : S \to \mathbb{R}$, $\alpha = (\alpha_1, E_1; \alpha_2, E_2; \ldots; \alpha_N, E_N)$ that satisfy the ranking $\alpha_1 > \alpha_2 > \ldots > \alpha_N$, the Choquet integral of α with respect to a capacity v is defined as:

$$\int \alpha dP = \sum_{n=1}^{N-1} (\alpha_n - \alpha_{n+1}) v \left(\bigcup_{m=1}^{n} E_m \right) + \alpha_N$$

This expression can be written as:

$$\int \alpha dP = \sum_{n=1}^{N} \alpha_n \left[v \left(\bigcup_{m=1}^{n} E_m \right) - v \left(\bigcup_{m=1}^{n-1} E_m \right) \right]$$

We next apply this definition of the *Choquet* integral to a *utility maximisation* problem, using our sequential decision task with three states of the world. This is the way that a *resolute* decision maker would approach the problem. We assume a state space $S = \{s_i, s_j, s_k\}$ which corresponds to the three different states that can be observed in the world, the set $X = \{x_i, x_j, x_k\}$ is the set of the corresponding consequences at each state of the world and v the subjective capacity of non-additive probabilities. We then need to specify 6 capacities, three capacities-one for each state $(v\{s_i\}, v\{s_j\}, v\{s_k\})$ and three for the joined capacities over states $v\{s_i \cup s_j\}$ for s_i, s_j, $v\{s_j \cup s_k\}$ for s_j, s_k and $v\{s_i \cup s_k\}$ for s_i, s_k .

In order to proceed, the model requires that we have a complete preference on F (Gilboa, 2010). All the functions from states to outcomes should be considered at least as conceivable acts. Thus, if there is a set of available acts E and constructs $S = X^E$ then F is defined as $F = X^S = X^{(X^A)}$ and we assume that the preference relationship $\succsim$ is a complete order on F. Firstly we need to rank the preferences of

[23] More precisely, the definition is based on the notion of preference for mixtures or differently hedge against ambiguity.

the decision maker from the best to the worst outcome, for example $x_k > x_j > x_i$

Summarizing the values on the following Table:

		s_i	s_j	s_k	$s_i \cup s_j$	$s_j \cup s_k$	$s_i \cup s_k$	S
v	0	v_i	v_j	v_k	$v\{s_i,s_j\}$	$v\{s_j,s_k\}$	$v\{s_i,s_k\}$	1
$u(x_s)$		$u(x_i)$	$u(x_j)$	$u(x_k)$				

For $x_i < x_j < x_k$ the Choquet integral is given by:

$$u(x_i)[1 - v(\{s_j,s_k\})] + u(x_j)[v(\{s_j,s_k\}) - v(\{s_k\})] + u(x_k)[v(\{s_k\})$$

or in a different way:

$$\int_S u(x_s)dv(s) = u(x_i) + (u(x_j) - u(x_i))v(\{s_j,s_k\}) + (u(x_k) - u(x_j))v(\{s_k\}) \quad (2.12)$$

and by rearranging:

$$\int_S u(x_s)dv(s) = u(x_i)[1 - v(\{s_j,s_k\})] + u(x_j)(v(\{s_j,s_k\}) - v(\{s_k\})) + u(x_k)v(\{s_k\})$$

The Choquet integral, can be interpreted as a *weighted* average of the expected utility, with weights that depend on the ranking of the different outcomes and sum up to 1. Let ω_s being the weight for state s. The weights $\omega_i,\omega_j,\omega_k$ for the ranking that we assumed, are defined in the following way:

$$\begin{aligned} \omega_k &= v_k \\ \omega_j &= v_{jk} - v_k \\ \omega_i &= 1 - vjk \end{aligned}$$

Defining the weights in such a way, allows us to write the Choquet integral as:

$$\int \alpha dP = \sum_{s=i}^{S} \omega_s u(x_s)$$

As in the case of MEU, the CEU model relaxes the *Independence* axiom. Instead, it substitutes this axiom with the one of *Comonotonic Independence*[24]. In order for the representation presented above to hold, it must be guaranteed that the acts

[24]Comonotonic Independence requires that $f \succ g \Rightarrow \alpha f + (1-\alpha)h \succ \alpha g + (1-\alpha)h$ only if f,g,h are pairwise comonotic. A pair of acts is comonotonic when there is no pair of states, s,s' such that $f(s) \succ f(s')$ and $g(s) \prec g(s')$.

under consideration are *comonotonic.* The optimal allocation for a *resolute* decision maker is obtained by the optimisation of the objective function that considers the problem as a static one. The conditional allocations will be the same as the initial ones. The objective is to maximise:

$$\omega_i u(z_i) + \omega_j u(z_j) + \omega_k u(z_k)$$

subject to the budget constraint $m = x_i + x_j + x_k$ and to the *ranking* that the outcomes must satisfy. This means that each ranking generates a different constrained optimisation problem which requires the incorporation of the Kuhn-Tucker conditions. In the case of three outcomes, the number of these additional constraints, is a function of the number possible combinations one may have. Three outcomes can be ordered in 3!=6 ways therefore, in order to completely characterise the solution, it is necessary to solve the problem for all the possible rankings and apply as acceptable allocation the one that maximises the utility of the agent. In Appendix A, we analytically present all the possible rankings, the respective weights that each ranking implies, as well as the analytical solutions for a CRRA utility function.

The decision process of a *naive* decision maker involves integration of sequential arrival of information and updating of the prior beliefs. Hence, an appropriate story is needed as to how this updating process takes place. The timing of the decision procedure for this type is the following: at time $t = 0$ she chooses the allocation that maximises the current utility (exactly in the same way that a resolute decision maker does). When the partial information becomes available, the capacities that characterise the decision maker's beliefs are updated and then a new optimisation problem is solved, based on her preferences and her updated beliefs. Let $\widetilde{w}$ be the updated weight. Since there are only two alternatives, only one weight is needed as one prospect is weighted by $\widetilde{w}$ and the other by $1 - \widetilde{w}$. Let for instance the available information to be that the state of the world is *not* j. The decision has to be made between assets i and k. Keeping the same notation as before and denoting $x_{i\neg j}, x_{k\neg j}$ the conditional allocations to i and k,the maximisation problem can be now written as:

$$\max_{x_{i\neg j}, x_{k\neg j}} \widetilde{wq} u(x_{i\neg j}) + (1 - \widetilde{w}) u(x_{k\neg j}) \tag{2.13}$$

$$\text{s.t. } x_{i\neg j} + x_{k\neg j} = m - x_j^* \tag{2.14}$$

The next step is to define how these weights are formed. In the literature, there have been suggested several different rules that are based either on the *statistical* approach of updating (Dempster (1967), Dempster (1968), Shafer (1976), Jaffray (2008)) or on the *decision theoretic* approach (Eichberger et al. (2007), Eichberger et al. (2010), Chateauneuf et al. (2011)).[25] In this chapter, we focus on the decision theoretic approach as we regard capacities as beliefs. Nevertheless, this approach requires the satisfaction of a set of axioms. The main result of Eichberger et al. (2007) is that the utility function over outcomes remains invariant to updating. Then, *consequentialism, state independence*[26] and *conditional certainty equivalent* consistency must hold when one wants to apply the available updating rules for CEU preferences.

The three most common rules that are used in order to update Choquet capacities are the *Optimistic*, the *Dempster-Shafer* and the *Generalised or Full Bayesian* update rule. We denote as $v(i|\neg j)$ the updated capacity for state i conditional on the information that the state of the world is not j. For an event $E \in \Sigma$ the conditional capacities for an event A are given by the following formulas:

Optimistic rule:

$$v_E^{OPT}(A) = \frac{v(A \cap E)}{v(E)}$$

Dempster Shafer rule:

$$v_E^{DS}(A) = \frac{v((A \cup E^c) - v(E^c)}{(1 - v(E^c))}$$

[25] The difference among the two lies on the way that the two methods approach the notion of capacity. In the case of the statistical approach the capacities represent objective, imprecise information and therefore, are generated by probabilities intervals. On the contrary, the decision theoretic approach considers capacities as subjective beliefs and is based on a series of axioms regarding the agent's preferences.

[26] The State Independence axiom entails that the ordinal ranking of outcomes remains unchanged no matter the event on which the preferences are being conditioned.

Generalised Bayesian Updating Rule:

$$v_E^{GB} = \frac{v(A \cap E)}{v(A \cap E) + 1 - v(A \cup E^c)}$$

which all collapse to the standard Bayesian rule in case that the capacity is additive. Then taking the case where we have a binary bet of the form xAy which yields outcome x on A and y on A^c with $x, y \in X$ and $x \succ y$. Thus it holds that $u(x) > u(y)$ and the CEU can be expressed as:

$$v(A)u(x) + (1 - v(A^c))u(y)$$

defining the conjugate capacity as $\tilde{v}(B) = 1 - v(B^c)$ the CEU can be written as:

$$V(xAy) = v(A)u(x) + \tilde{v}(A^c)u(y)$$

Then after obtaining information that the event is E, the updated CEU becomes:

$$V_E(xAy) = v_E(A)u(x) + \tilde{v}_E(A^c)u(y)$$

and applying each updating rule we obtain:
Optimistic:

$$v_E^{OPT}(xAy) = \frac{v(A \cap E)}{v(E)}u(x) + \frac{v(E) - v(A \cap E)}{v(E)}u(y)$$

Dempster-Shafer rule:

$$v_E^{DS}(xAy) = \frac{v((A \cup E^c) - v(E^c)}{(1 - v(E^c))}(x) + \frac{1 - v((A \cup E^c)}{(1 - v(E^c))}u(y)$$

Generalised Bayesian Updating Rule:

$$v_E^{GB} = \frac{v(A \cap E)}{v(A \cap E) + 1 - v(A \cup E^c)}u(x) + \frac{1 - v(A \cup E^c)}{v(A \cap E) + 1 - v(A \cup E^c)}u(y)$$

As Eichberger et al. (2010) remark, the Optimistic updating rule, weighs the good outcome following the rule of conditional probabilities, while the bad outcome is determined by the complementary decision weight. On the contrary, the Dempster-Shafer rule puts more weight on the bad outcome based on the law of conditional probabilities and the weight to the good outcome is the complementary weight. Finally, the Generalised Bayesian Updating rule, adjusts both weights in such a way that they are symmetric for both the good and the bad outcome. All of the rules that have been axiomatized, share the same property, that the utility index applied to represent the updated preferences remains unchanged. In Appendix A we present how these rules are applied to our problem taking into consideration the relevant rankings.

Table 2.7 shows the simulated decisions of a Naive decision maker with CEU preferences, based on 25 problems out of total 60 that were presented to the subjects during the experiment, assuming that the information received is that the state of the world is *not k*. We show what the allocations are for assets i and j and then the respective allocation conditioned to the event. In the first stage ($t = 0$), the decisions (x_i, x_j) coincide independently of which updating rule is applied as the decision maker is solving the one-stage static problem. They also coincide with what a Resolute type would do. In the next columns ($t = 1$), the agent uses the available information and updates the beliefs accordingly, leading to the respective allocations $(x_{i\neg k}, x_{j\neg k})$. The set of parameters according to which the simulation was carried out is presented in Table 2.7. No stochastic component is assumed and the utility is represented by a constant relative risk aversion (CRRA) function.

Table 2.6 Parameter Values for CEU Simulation

Parameter		**Value**
v_i	=	0.180
v_j	=	0.270
v_k	=	0.400
v_{ij}	=	0.520
v_{ik}	=	0.790
v_{jk}	=	0.680
r	=	0.300

From Table 2.7 it is possible to see that different updating rules generate different optimal allocations and thus, the experimental design enables us to identify which rule is being used.

Table 2.7 Optimal Allocations for CEU Naive

					$t=0$				$t=1$			
							OPT		DSU		GBU	
Problem	**Income**	e_i	e_j	e_k	x_i	x_j	$x_{i\neg k}$	$x_{j\neg k}$	$x_{i\neg k}$	$x_{j\neg k}$	$x_{i\neg k}$	$x_{j\neg k}$
1	10	1.3	1.5	1.4	1.62	4.04	2.59	3.08	1.59	4.08	2.25	3.42
2	20	0.6	0.8	0.7	6.21	4.66	4.80	6.07	2.91	7.96	4.15	6.72
3	30	0.3	0.5	0.4	4.39	11.38	6.60	9.17	3.93	11.84	5.67	10.10
4	40	0.2	0.4	0.5	5.01	16.34	8.53	12.82	5.01	16.34	7.30	14.06
5	50	0.2	0.3	0.3	7.28	20.97	12.14	16.12	7.28	20.97	10.45	17.80
6	60	0.2	0.3	0.3	8.74	25.17	14.56	19.34	8.74	25.17	12.54	21.36
7	70	0.2	0.2	0.2	11.80	28.57	19.08	21.30	11.80	28.57	16.60	23.77
8	80	0.1	0.2	0.2	10.47	34.13	17.83	26.78	10.47	34.13	15.24	29.37
9	90	0.1	0.2	0.2	11.78	38.40	20.06	30.13	11.78	38.40	17.14	33.04
10	100	0.1	0.2	0.2	13.09	42.67	22.28	33.48	13.09	42.67	19.05	36.71
11	10	1.4	1.5	1.3	1.69	3.86	2.58	2.97	1.59	3.96	2.24	3.31
12	20	0.7	0.8	0.6	3.40	7.12	4.82	5.69	2.95	7.56	4.18	6.33
13	30	0.4	0.5	0.3	5.17	9.31	6.50	7.98	3.95	10.53	5.62	8.86
14	40	0.5	0.4	0.2	8.31	10.92	10.68	8.55	6.01	13.22	10.68	8.55
15	50	0.3	0.3	0.2	9.12	16.35	12.04	13.44	7.45	18.03	10.47	15.00
16	60	0.3	0.3	0.2	10.95	19.62	14.44	16.12	8.93	21.63	12.57	18.00
17	70	0.2	0.2	0.2	11.80	28.57	19.08	21.30	11.80	28.57	16.60	23.77
18	80	0.2	0.2	0.1	15.55	22.26	17.87	19.94	11.05	26.76	15.55	22.26
19	90	0.2	0.2	0.1	17.49	25.05	20.10	22.44	12.43	30.10	17.49	25.05
20	100	0.2	0.2	0.1	19.43	27.83	22.33	24.93	13.81	33.45	19.43	27.83
21	15	1.4	1.5	1.3	2.53	5.79	3.87	4.45	2.38	5.94	3.36	4.96
22	25	0.7	0.8	0.6	4.25	8.89	6.02	7.12	3.69	9.45	5.22	7.92
23	35	0.4	0.5	0.3	6.04	10.86	7.58	9.31	4.61	12.29	6.56	10.34
24	45	0.5	0.4	0.2	9.35	12.28	12.02	9.61	6.76	14.87	12.02	9.61
25	55	0.3	0.3	0.2	10.03	17.99	13.24	14.78	8.19	19.83	11.52	16.50

The decision task for a *sophisticated* decision maker, requires a solution using backward induction. As was analytically presented in section 2.2.2, this requires a kind of connection between conditional and unconditional expectations. This means, that initially, the decision maker should solve the conditional problems based on the information that she may receive and applying the respective updating rules as shown above, and then solve the global unconditional problem, in order to determine the unconditional allocations. In this case, the decision maker faces three conditional states that each has two outcomes. This means that when solving for the global unconditional problem, the agent should consider all the six conditional outcomes and their perspective ranking. With 6 outcomes, there are 6!=720 possible different rankings that one should consider, applying at each the correct weights. In Appendix A, we present the analytical solutions for the problem, as well as the steps of the algorithm that we developed in order to account for the 720 rankings.

2.2.5 Source Choquet Expected Utility (SCEU)

The source method is based on the work of Tversky and Kahneman (1992) on the different sources of ambiguity and has been recently developed and used in the analysis of experimental data by Abdellaoui et al. (2011) and Kothiyal et al. (2014). The main idea of this approach is that different sources of uncertainty can be treated in a different way. As Kothiyal et al. (2014) indicate, the use of these source functions allows the mapping of choice-based probabilities into willingness to bet. As a result, decision making under ambiguity is described by the use of three components, the utility over outcomes that represents tastes, the choice-based probabilities that reflect beliefs and the source functions that capture diversion from Expected Utility theory (Allais and Ellsberg type behaviour, home bias and ambiguity aversion). As a definition, a source of uncertainty concerns a group of events that is generated by a specific mechanism of uncertainty. In our framework, it is reasonable to assume that the Bingo Blower is a source of ambiguity, as the various events and their respective representation are created in a particular way. The sources are assumed to be algebras which means that they contain the universal event (happens with probability equal to 1), the empty event, the complement of each of the elements and the union for each pair of the elements. A source function is defined for a specific source of ambiguity and a specific treatment. That is to say, a subject i in a given treatment, holds subjective beliefs for the various states of the world that can be represented by a vector $\mathcal{P}$ over $\mathcal{S}$. Then, there exists a source function, which is a mapping in the interval [0,1] such that $w(0) = 0$, $w(1) = 1$ and $w(E) = w(P(E))$. The objective of this function is to transform the subjective probabilities. Using the family[27] that Prelec (1998) suggests, the function can be written as:

$$w(p) = \exp(-(-\ln(p))^{\alpha}) \tag{2.15}$$

The relationship above can be described with the help of a graph of the source function. The x-axis represents the probability p while on the y-axis the weighted-transformed probabilities are represented. When the coefficient α is equal to 1, there are no weighted probabilities and the beliefs satisfy the assumption of probabilistic sophistication. The source function can be characterised as linear and its shape coincides with the 45° line. This case is illustrated in Figure 2.4 on the top left graph. Then, depending on the value of α there may be either an S-shaped source function in the case that α is greater than one or an inverse S-shaped source function in the case that α is less than one. Similarly the weighting function for $\alpha = 0.5$, $\alpha = 1.5$ and $\alpha = 0.85$ are illustrated.

The intuition of characterizing a weighting function as an inverse S function is that people usually tend to overweight large probabilities and to underweight those that are larger and are associated with larger outcomes. Evidence on this can be found on Abdellaoui (2000) and Tversky and Kahneman (1992). Using this methodology there are some advantages and at the same time some drawbacks. As is already mentioned, the positive aspect of this methodology is that the number of parameters to estimate is significantly reduced. In the case of CEU preferences with three outcomes, there are six parameters for the capacities that represent beliefs, plus the coefficient of risk aversion r and the precision parameter s for the stochastic specification, requiring 8 parameters in total to be jointly estimated. Using the source method, these parameters are reduced to five, as now

[27] As the objective is to minimise the number of parameters to estimate, we follow Kothiyal et al. (2014) and apply the one-parameter family. In the same study, they discuss the alternative specifications (either with one or two parameters) that can be used.

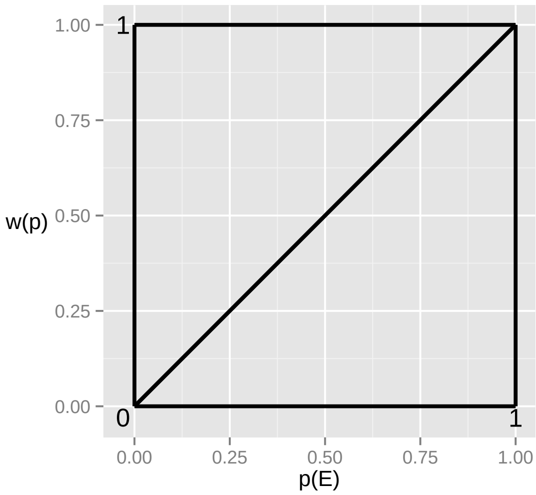

(a) Linear Function, $\alpha = 1$

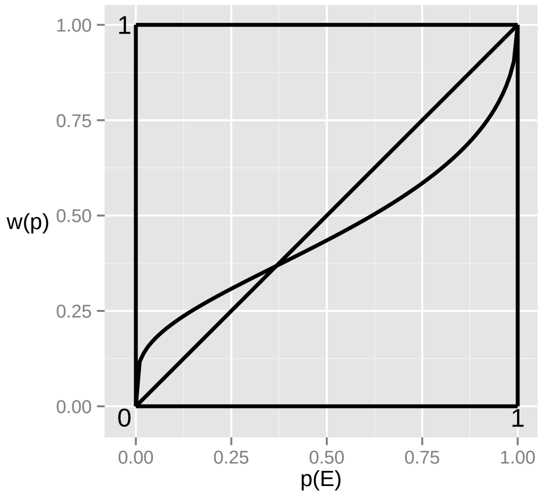

(b) Inverse-S Shape Function, $\alpha = 0.5$

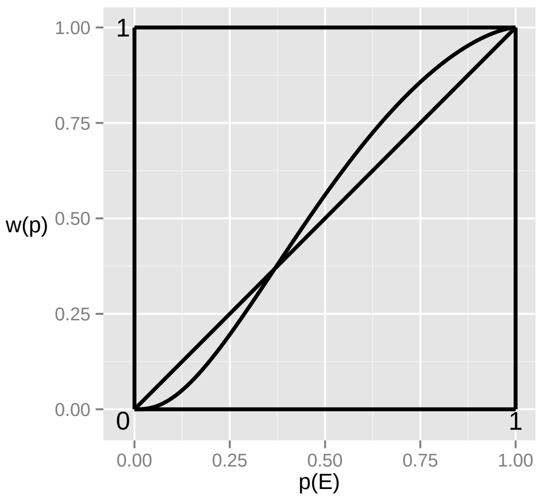

(c) S Shape Function, $\alpha = 1.5$

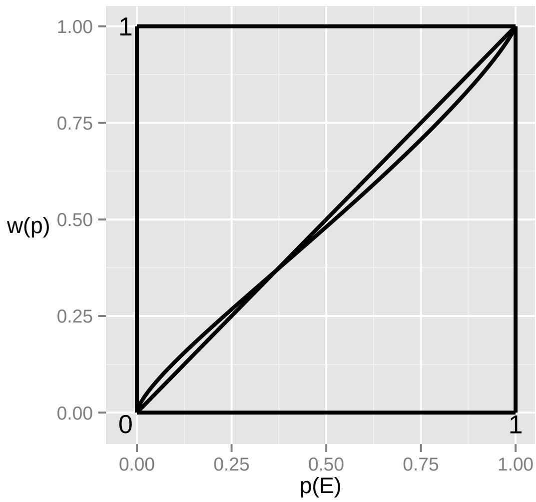

(d) Inverse S Shape Function $\alpha = 0.85$

Figure 2.4 Weighting of Probabilities

there are two probabilities to estimate (p_i, p_j and p_k is simply the residual), the parameter of the weighting function α and again the risk aversion coefficient and precision. In addition, on the positive side is that this kind of weighing offers a measure of ambiguity aversion, something that was previously measured only by the non-additivity of the capacities.

On the other hand, this specification significantly constraints one of the basic objectives of this research. As one of the scopes is to test how ambiguous beliefs are updated, the source function method assumes probabilistic sophistication, which does not allow the representation of beliefs using capacities and consequently all the alternative rules that have been suggested for the CEU model, are not possible to apply in this case. As a result, we need to make the assumption that upon the arrival of new information, the prior beliefs are updated based on the Bayesian rule. Nevertheless, it allows to test behaviour in this sequential framework with the only difference compared to the CEU being the lack of capacities and thus, different updating rules. As updating rules and behaviour (types) are interconnected, adapting this specification allows one to focus purely on the different types without considering the updating process. This kind of modelling belongs to the class of the Rank Dependent Expected Utility models, which means that as in the case of the CEU the ranking of the different outcomes $e_i x_i$ plays a significant role.

Again, for this model, we examine three different types, the resolute, the naive and the sophisticated. The *resolute* decision maker, anticipates only three possible states of the world and holds subjective beliefs $\{p_i, p_j, 1 - p_i - p_j\}$. Assuming the following ranking of the outcomes $e_i x_i > e_j x_j > e_k x_k$ and the weighting function A.61, the weights $w(.)$ attached to the outcomes are given by:

$$\omega_i = w(p_i) \tag{2.16}$$

$$\omega_j = w(p_i \cup p_j) - w(p_i) = w(p_i + p_j) - w(p_i) \tag{2.17}$$

$$\omega_k = 1 - w(p_i \cup p_j) = 1 - w(p_i + p_j) \tag{2.18}$$

$$\tag{2.19}$$

All the appropriate weights for all possible rankings are presented in Table A.7 of Appendix A. In addition, as the specification is quite similar to the CEU case, with the only difference being the way that the capacities are handled, the maximisation problem, as well as the solutions for the optimal allocations, are exactly the same as in the CEU case, adjusted for the appropriate weights (see Appendix A).

As already discussed, the *naive* decision maker will choose at the first stage exactly in the same way as a resolute decision making but is doing this unintentionally. In the second stage, and always assuming that the beliefs of the decision maker are weighted before making a decision, the choices at the second stage will deviate from what was chosen at $t = 0$. The nature of the problem to be solved at the second stage remains the same as before. So for instance, if the event E is that the ball is not i ($p(E) = p(s_j \cup s_k)$) then the optimisation task reduces to:

$$\max_{x_{j\neg i}, x_{k\neg i}} w'_j u(z_j) + w'_k u(z_k) \tag{2.20}$$

$$\text{s.t. } e_j x_{j\neg i} - e_k x_{k\neg i} \leq 0 \tag{2.21}$$

$$x_{j\neg i} + x_{k\neg i} = x^*_j + x^*_k \tag{2.22}$$

where w'_j, w'_k are the updated weights for j and k respectively. What remains now is to determine the way that the weights are updated. Since the subjective beliefs are additive, we are still in position to apply the Bayesian rule to update these beliefs. Then the updated subjective beliefs, are accordingly weighted by the weighting function $w(.)$ (Equation A.61). The conditional beliefs for states j and k on E are respectively given by:

$$p(s_j|E) = \frac{p(s_j)}{p(s_j \bigcup s_k)} = \frac{p(s_j)}{p(s_j) + p(s_k)} \tag{2.23}$$

$$p(s_k|E) = \frac{p(s_k)}{p(s_j \bigcup s_k)} = \frac{p(s_k)}{p(s_j) + p(s_k)} \tag{2.24}$$

Then, the relevant rankings[28] are taken into consideration in order to complete the weight matrix.

Finally, the *sophisticated* decision maker, backwardly solves the problem and attach the appropriate weights where needed in a similar manner as in the CEU case (details on Appendix A).

2.2.6 Specifications and Respective Axioms

It has become clear by now, that when one deviates from the SEU model, inevitably either the axiom of dynamic consistency or that of consequentialism is violated. In this chapter, we have presented 14 possible specifications (combinations of preference functionals, updating rules and types) that explicitly make some as-

[28] As in the CEU case, for each conditional state, three rankings are taken into consideration namely the rankings $z_j > z_k, z_j < z_k$ and $z_j = z_k$, where z is the payoff at each state.

sumptions, regarding which axiom is satisfied, topic that has been extensively discussed in chapter 1. Whenever updating is taking place (naive and sophisticated type) and appropriate rules are applied, consequentialism in assumed to hold, or stating differently, the choices are not dynamically consistent. Nonetheless, when the type is resolute, dynamic consistency is violated[29]. Doing such a classification, allows to consider case where, even though the decision maker is non-EU, she does not fail to satisfy dynamic consistency. In Table 2.8, the combinations of the specifications along with the respective axioms, are presented.

Table 2.8 Specifications and Respective Axioms

	Preferences	**Type**	**Updating Rule**	**Axioms satisfied**
1	SEU	Resolute	-	DC/C
2	MEU	Resolute	-	DC
3	MEU	Naive	GBU/MLU	C
4	MEU	Sophisticated	GBU/MLU	C
5	CEU	Resolute	-	DC
6	CEU	Naive	GB	C
7	CEU	Naive	DS	C
8	CEU	Naive	OPT	C
9	CEU	Sophisticated	GB	C
10	CEU	Sophisticated	DS	C
11	CEU	Sophisticated	OPT	C
12	SCEU	Resolute	-	DC
13	SCEU	Naive	BAYESIAN	C
14	SCEU	Sophisticated	BAYESIAN	C

C stands for *Consequentialism* and DC for *Dynamic Consistency*

2.3 The Experimental Procedure

The experimental design is inspired by the *sequential* Ellsberg problem that was presented in chapter 1, but is substantially different in nature, regarding the representation of ambiguity, as well as the decision problem. In our experimental protocol, subjects had a rough idea of the composition of the urn (there was not complete ignorance regarding the probabilities of some states, a fact that reduces

[29]In this case, the dynamic consistent updating rules (Klibanoff and Hanany (2007), Hanany and Klibanoff (2009)) or the methods that the recursive models propose (Epstein and Schneider (2003)) can be explained by the resolute type-but cannot be identified.

the level of suspicion that a standard Ellsberg urn may entail[30]) and in addition, they were required to realise allocations of experimental income to the possible events instead of choosing binary bets. We use a *within-subject* design as we are focusing on individual choice. The decision task in the experiment is quite simple: using a Bingo Blower to represent ambiguity in the lab, the subjects were presented with a series of allocation problems, with the objective being to allocate the experimental income to the three assets based on the probabilities that each asset has to be extracted from the Bingo Blower. The problems are constrained to allocating an experimental income m to three *Arrow* security assets ($x_s, s \in \{i, j, k\}$) where an Arrow security is defined as the asset that pays 1 monetary unit if the state of the world is s and zero otherwise.[31]

The ambiguous environment was created with the use of a *Bingo Blower*. The Bingo Blower consists of a transparent box that contains table tennis balls. At the bottom of the box, there is a motor that generates a stream of air, which makes the balls to continuously move inside the box. The advantage of this device, is that when the number of the balls is sufficiently high, one is not able to count the total number of the balls, or the number of individual colours. What is possible to do, is to distinguish that there is at least one ball of each colour (minimum probability) and to obtain a rough idea of the maximum number of the balls (maximum probability) preserving always some ambiguity.[32] In other words, while there exist objective probabilities (known to the experimenter), the subjects are not able to precisely construct an objective probability distribution. Inside the Bingo Blower there were balls of three different colours Blue, Yellow and Pink, to represent the three different Arrow assets. A similar Bingo cage has been used by Andersen et al. (2012) and the Bingo Blower has been used by Hey et al. (2010) and Hey and

[30] A standard practice in laboratory experiments that include some form of Ellsberg urns, is to provide subjects with very vague information of the kind "there are 100 Black and Red balls but their ratio is unknown" or "there may be 0-100 Black balls". This type of information creates suspicion that the experimenter is trying to minimise the payments, as she or he does not explicitly describe how the urn was composed.

[31] In this case it is assumed that the exchange rates are equal to 1 and the experimental income has been normalised to 1. This means that the payoffs above correspond to the case where the entity of the income is allocated to the asset s. In the experiment, there were different exchange rates in each problem and the payoff at state s was the product between the proportion x_s of the total income allocated to the asset s multiplied by the respective exchange rate e_s.

[32] Roughly, when $n > 10$ the environment becomes ambiguous enough. When $n < 10$, it is possible that subjects will be able to count the exact number of balls, transforming the problem to a risky one.

Pace (2014), all in *static* choice problems. The actual synthesis of the Bingo Blower during the experiment is shown in Table 2.9.

Table 2.9 Composition of the Bingo Blower

	Blue	Yellow	Pink	TOTAL
Number of balls	4	12	6	20
Percentage	20%	50%	30%	1

The decision task can be described with the use of a three-period model. At $t = 0$ the subjects received an experimental income m and they were informed on the exchange rates (e_i, e_j, e_k) between experimental income and British sterling that each asset yields. Based on these exchange rates, they were asked to make an allocation to the three assets, knowing that only one will be the *winning* asset and as such, the only one that provides a payoff equal to the amount allocated to this asset (colour), transformed to monetary value. But before learning the actual state of the world, the subjects were receiving some relevant information that would help them to re-evaluate their choices, if needed, and were given the opportunity to change their initial choice. The subjects were asked to imagine[33] the following scenario:

> A ball is drawn from the Bingo Blower. We will not reveal you the actual colour of the ball but we will let you know what the colour of the ball that we have in our hands *is not*. Then, you will lose all the income that you allocated to this "not" ball, since it is not the *winning* colour for sure.

This means, that at $t = 1$ the subjects received partial information in the form of a filtration that presents the complementary event (e.g. $\mathcal{F}_1 = \{\{Blue, Pink\}, \{Yellow\}\}$ which provides the information whether the ball is *Yellow* or not-subjects were never received the filtration saying that an event happened with certainty). At that time, the subjects were asked to allocate the remaining available income (initial income-amount allocated to the "not" colour) to the two possible states of the world (e.g. to Blue and Pink) and at $t = 2$ all relevant uncertainty is resolved by revealing what the colour of the ball actually is. There were 60 similar problems[34],

[33] In the experiment, we asked subjects to imagine that a ball is drawn, instead of actually drawing, in order to avoid any learning effects. This is discussed in section 2.3.

[34] As the nature of the problem we investigate is quite complex, it is inevitable to demand a rich

with different amounts of experimental income and exchange rates) and the subjects were paid using an *incentive compatible* mechanism[35] by picking at random one of the 60 problems and physically playing it for real. The computer retrieved the allocations of each subject as well as the information that they were given at this specific problem. Then, they activate the mechanism that extracts balls from the Bingo Blower in order to randomly define the winning colour.[36]. Then, they were paid in cash and in private after the end of the experiment. The problems appeared in a random order, different for each subject to eliminate possible *order effects*.

2.3.1 The Problems

Before running the actual experiment we carried out a pilot study (it was conducted in two sessions with 6 participants each) to test the set of problems and the experimental software. Then, a series of extensive simulations was performed in order to ensure that the problem set that we have is useful to identify the different types. The set of questions was chosen in such a way that it could guarantee two criteria. First, the expected payoff for a *risk neutral* person would always be the same. A risk-neutral person devotes everything to the colour s which implies the highest product between the probability of event s and the exchange rate e_s. If there were no exchange rates between experimental income and money, the optimal choice would be to allocate everything to the asset with the highest probability to happen, thus to allocate everything to Yellow. In order to preserve this property, we chose the problems in such a way that a risk neutral person will always choose to allocate everything to the most probable outcome, Yellow. The second criterion on choosing the problems for the experiment was the performance of these specific problems in the testing of the estimation programmes. Having as an objective the development of robust estimation codes, we simulated and estimated experimental data with many different sets of problems (Monte Carlo simulation). The set that we applied in the experiment was the one that provided the most robust

dataset in order to obtain as accurate estimates as possible.

[35]For a support of this mechanism see Wakker (2007).

[36]The subjects were keeping drawing balls, with replacement, if the ball extracted was the same colour with the *not* ball.

estimators (maximum likelihood estimation). Table 2.10 shows the set of problems that we actually used in the experiment.

For each of the subjects, we recorded data on the decisions they made in this three-period task. There were 60 questions, each consisted of a two-stage decision, giving in total 120 observations per subject and pooling the data for the whole population of the experiment, we have $58 \times 120 = 6960$ observations. For the conditional states, the available data are roughly $1/3$ of the total dataset for each conditional state.[37] The structure of the experiment is such that allows us to proceed in both a subject level analysis and in aggregate by setting up a mixture model[38]. In this chapter we follow a subject by subject level analysis, as this provides the flexibility to fit a wide collection of different combinations[39] (preference functionals, types and updating rules).

[37] The experimental software was programmed in such a way, that each conditional state was appearing with probability equal to $1/3$.

[38] This is left for future work.

[39] In this chapter we specify and fit 14 different preference functionals for each of the subjects. See Table 2.12.

Table 2.10 Problems of the Experiment

Problem	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**
m	10	20	30	40	50	60	70	80	90	100	10	20	30	40	50	60	70	80	90	100
e_i	1.3	0.6	0.3	0.2	0.2	0.2	0.2	0.1	0.1	0.1	1.4	0.7	0.4	0.5	0.3	0.3	0.2	0.2	0.2	0.2
e_j	1.5	0.8	0.5	0.4	0.3	0.3	0.2	0.2	0.2	0.2	1.5	0.8	0.5	0.4	0.3	0.3	0.2	0.2	0.2	0.2
e_k	1.4	0.7	0.4	0.5	0.3	0.3	0.2	0.2	0.2	0.2	1.3	0.6	0.3	0.2	0.2	0.2	0.2	0.1	0.1	0.1

Problem	**21**	**22**	**23**	**24**	**25**	**26**	**27**	**28**	**29**	**30**	**31**	**32**	**33**	**34**	**35**	**36**	**37**	**38**	**39**	**40**
m	15	25	35	45	55	65	75	85	90	110	10	20	30	40	50	60	70	80	90	100
e_i	1.4	0.7	0.4	0.5	0.3	0.3	0.2	0.2	0.2	0.2	1.8	0.9	0.6	0.5	0.4	0.3	0.3	0.2	0.2	0.2
e_j	1.5	0.8	0.5	0.4	0.3	0.3	0.2	0.2	0.2	0.2	1.5	0.8	0.5	0.4	0.3	0.3	0.2	0.2	0.2	0.2
e_k	1.3	0.6	0.3	0.2	0.2	0.2	0.2	0.1	0.1	0.1	1.6	0.8	0.5	0.4	0.3	0.3	0.2	0.2	0.2	0.2

Problem	**41**	**42**	**43**	**44**	**45**	**46**	**47**	**48**	**49**	**50**	**51**	**52**	**53**	**54**	**55**	**56**	**57**	**58**	**59**	**60**
m	9	14	21	29	35	39	48	54	67	78	10	20	30	40	50	60	70	80	90	100
e_i	1.8	1.2	0.8	0.6	0.5	0.5	0.4	0.4	0.3	0.3	1.5	0.8	0.5	0.4	0.3	0.3	0.2	0.2	0.2	0.2
e_j	1.7	1.1	0.7	0.5	0.4	0.4	0.3	0.3	0.2	0.2	1.5	0.8	0.5	0.4	0.3	0.3	0.2	0.2	0.2	0.2
e_k	1.6	1	0.6	0.4	0.3	0.3	0.2	0.2	0.1	0.1	1.5	0.8	0.5	0.4	0.3	0.3	0.2	0.2	0.2	0.2

m is the experimental income, e_i is the exchange rate for *Blue*, e_j for *Yellow* and e_k for *Pink*.

2.3.2 Procedure and Administration

The experiment was conducted at the Centre for Experimental Economics (EXEC) at the University of York between May and June 2013. The subjects were recruited from a standard student experimental population using the ORSEE system (Greiner, 2004). The majority of subjects were Bachelor students from several different majors and 52% were females. There were three sessions (24,17,17). The experiment lasted for less than one hour. The subjects were paid privately in cash directly after the end of the experiment. The average payment was £11.16 including a *show-up* fee of £3. The maximum payment was £25.5. Upon arrival, the subjects were randomly assigned to a computer terminal.

The experiment was computerised and the experimental interface was developed in *Python*.[40] Each problem that the subjects were required to answer had two stages. Figures 2.5 and 2.6 show screenshots from the experimental framework for stages 1 and 2 respectively.

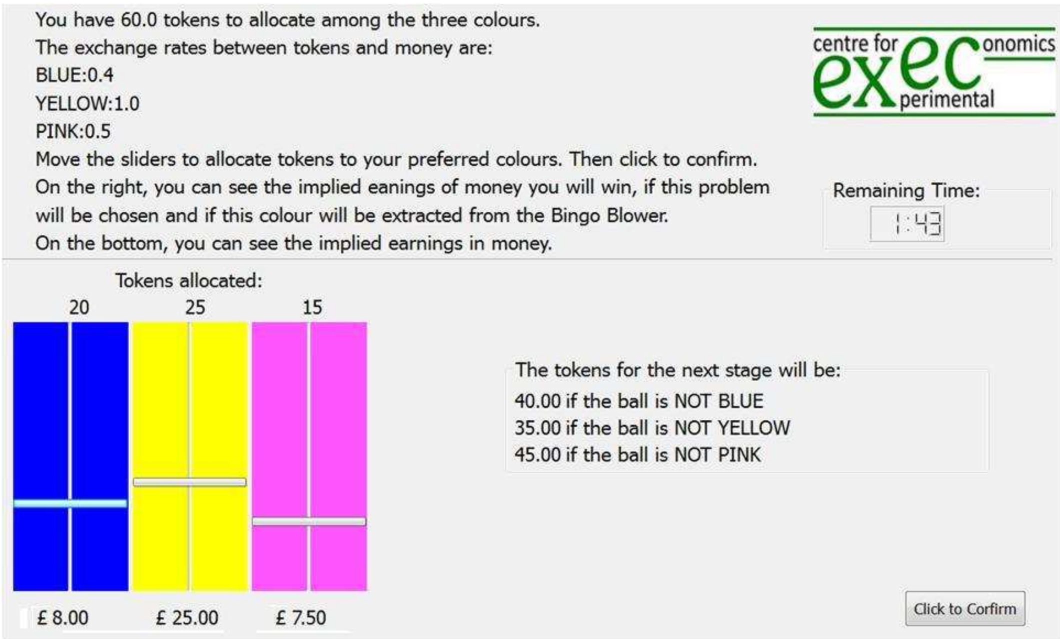

Figure 2.5 Experimental Software - Stage 1

In the first stage screen, the subjects could see three sliders, one for each respective colour, and the data for each problem: their total income to allocate and the exchange rates of the three colours. All the sliders were programmed to be

[40]Python Software Foundation. Python Language Reference, version 2.7. Available at http://www.python.org. The software is available upon request.

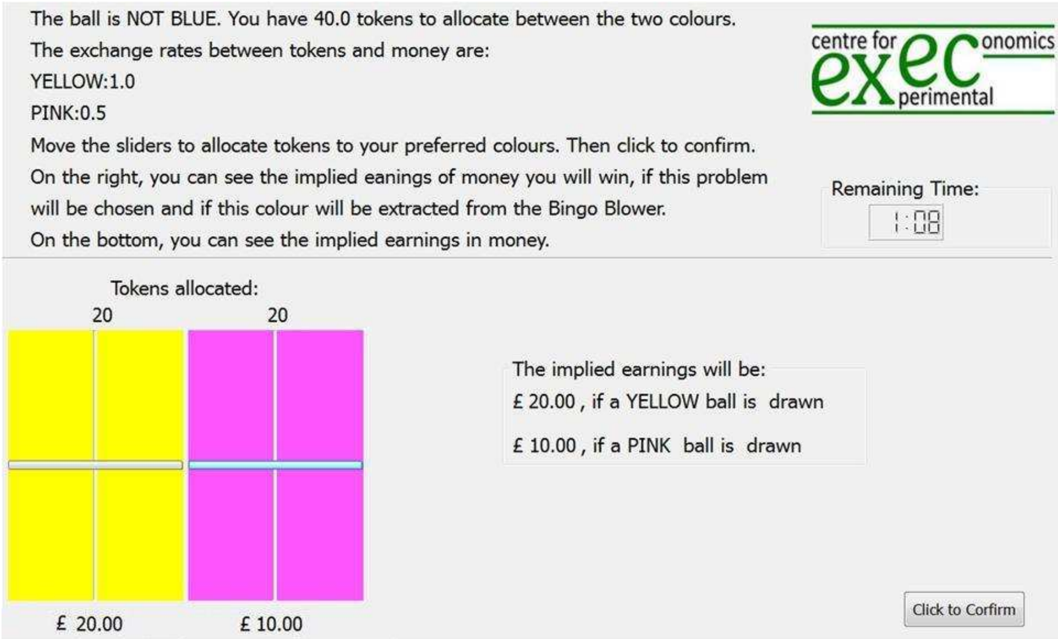

Figure 2.6 Experimental Software - Stage 2

connected with each other, so at any time, the allocation of the subjects was constrained to be equal to the total experimental income (there was no possibility to short-sell or to save for the next round). In addition, in the screen they could see how much of the income they have allocated to each colour and the respective potential income that they will have in the second stage, depending on the real state of the world. Also, there was a timer, with the minimum time to respond set to 10 seconds and the maximum to 90 seconds. Pushing the "next" button, the software was programmed to reveal some partial information (not i, not j, not k,) based on a uniform distribution. The latter was applied in order to avoid learning effects that would transform the problem to a risky choice one. In the second stage, there were only two sliders available for remaining states of the word, along with all the relevant information.

All subjects were given both written instructions (available in Appendix C) and a slide-show presentation. After reading the instructions, subjects were able to go through the slide-show presentation which was available at each terminal and contained simplified instructions and examples, and navigate it at their own pace. Then, they were free to go near the Bingo Blower and observe the composition of the three colours where they could create some priors on the proportions. During the experiment a live image of the bingo blower was projected through two large screens in the lab and in addition, during the experiment the subjects were free to walk around and observe the Bingo Blower whenever they wanted to. The choices were recorded in discrete values (integer values) in the range [0,m] with integer steps.

2.4 The Econometric Specification

In order to be in position to estimate the parameters for the various models, we need to make a series of assumptions regarding the utility function and the stochastic part of the decision making process. By specifying the decisions of the subjects using parametric forms of the models, it is possible to *jointly* estimate all the parameters of interest. This methodology differs from the approach of gradually estimating some of the parameters and substituting the estimated values so as to derive the remainder of the values (e.g. Baillon et al. (2013)) or using calibration methods and applying values for specific parameters based on what other studies have found.[41]

2.4.1 Utility Function

We assume that the preferences of the subjects are represented by a Constant Relative Risk Aversion (CRRA) utility function that has the following form:

$$u(x) = \begin{cases} x^r & \text{if } r > 0 \\ \ln(x) & \text{if } r = 0 \\ -x^r & \text{if } r < 0 \end{cases}$$

where x is the respective income (payoff) and r is the parameter of risk aversion. We choose to use this following specification for two reasons. On the one hand, it is generally acceptable that in empirical work this function provides the best fit (Stott (2006), Wakker (2008)) and specifically in our case, the CRRA function indeed provided better fit.[42] On the other hand the CRRA function has a very

[41] While this methodology is quite efficient as it reduces the number of parameters to estimate, it is restricting as it does not allow for full heterogeneity among subjects.

[42] We have fitted the models using the CARA and the ExpoPower function. While the CARA utility function has the nice property of allowing for boundary allocations, the differences with the CRRA function were trivial. The ExpoPower on the other hand, has nice properties but it faces two important limitations. Firstly, it adds one extra parameter that needs to be estimated. Then, due

convenient property that it does not allow for boundary portfolios (zero allocations or allocations to only one set). Practically, this means that with the use of Power function, it is possible to derive elegant, *closed-form* formulas for the optimal allocations.

If one wants to assume an exponential utility function (CARA), then there is a positive aspect, where allocations on the bounds can be explained behaviourally, but modelling utility in such a way creates a series of problems. Firstly, the unconstrained optimal allocations are not constrained to the feasible interval defined by the experimental design [0,m], but can be zero (boundary) or even negative (short-selling-which was not allowed during the experiment). As a consequence, the form of the optimal solution changes as now it is necessary to check all he possible combinations that may appear when the strictly positive allocation to all the assets ($x_i^* > 0, x_j^* > 0, x_k^* > 0$) is not satisfied. We have developed an algorithm (details are provided below) that solves for the optimal allocations by taking into consideration all the possible combinations and the respective constraints, but assuming a CARA utility function creates some additional difficulties. As explained later in the stochastic specification, when the optimal allocation is zero and an actual allocation is observed to be strictly positive, the assumed distribution of the random variables degenerates and it becomes impossible to estimate the preference functionals. In order to overcome similar issues, we suggest different ways of modelling stochastic choice (this is extensively discussed in chapter 3). These alternatives solve the degeneration issues but come at a cost, that of an extra parameter to estimate. As the preference functionals that we assume contain already a relatively high number of parameters (the minimum number is in the case of SEU with 4 parameters and the maximum the CEU case with 8) we decided to use the CRRA specification with the simplest error story that we can suggest. In addition, observing the data, there are very few subjects that have chosen boundary portfolios, making the use of a CRRA function appropriate. The properties of the CRRA utility function are summarised in Table 2.11.

One of the constraints in this experimental design, is that it does not allow the modelling of a risk loving person who would be willing to short-sell and allocate

to its functional form, it is not possible to obtain a closed form solution. Consequently, numerical methods should be applied that both increase the computational time and decrease the precision. The latter was not applied to the more complicated models (e.g. CEU).

Table 2.11 Risk Attitudes

$u(x)$	$\frac{\partial u}{\partial x}$	$\frac{\partial^2 u}{\partial x^2}$	**Attitude**	**r**
x^r	rx^{r-1}	$(r-1)rx^{r-2}$	Averse	$r<1$
$log(x)$	$1/x$	$-1/x^2$	Neutral	$r=1$
$-x^r$	$-rx^{r-1}$	$-(r-1)rx^{r-2}$	Loving	$r>1$

negative amounts to one of the assets.

The estimation of the parameters is performed by using Maximum Likelihood Estimation (MLE) techniques. The estimation program was written in the *R* programming language for statistical computing[43].

2.4.2 The Stochastic Specification

In order to complete our story of how the decision process takes place, we need to make some assumptions about the stochastic specification. Incorporating stochastic choice is of paramount importance as on the one hand it makes the estimation of the several specifications feasible and on the other hand, it renders the modelling of decision making more realistic. Hey (2014), explains that the errors in the decision making are either due to randomness in the preferences or due to errors that the subjects make when they choose. The existence of noise can be explained by the fact that the subjects make mistakes when they make decisions. This can be due to several factors, such as difficulties to understand the problems, badly designed software or low incentives. The two former can be corrected by increasing the number of questions and repetitions (learning is a factor that reduces noise). Another model of noise is the *random preferences* model where the subjects are assumed to have a set of different preference functionals, not only one, when they make decisions. The error story that will be adopted seems to be quite significant as this will define the most appropriate econometric method to be used. Wilcox (2011), names the parameter for the error as tremble probability. This is the probability that there is some randomness in the observed choice due to attention lapses or responding mistakes.

In our experiment, for each question, we can solve for the optimal allocation vector x^*, according to the preferences and the utility representation. Since we focus on continuous random variables, it is easy to think it terms of proportions

[43]The *R* Manuals, version 3.0.2. Available at: http://www.r-project.org/

over the total available income x_i/m. Due to the constraint $0 \leq x_i \leq m$, this amount is always constrained by the interval [0,1]. Thus, the most natural way to model the stochastic process is by assuming that x_i/m follows a *Beta* distribution. The *Beta* distribution is a family of continuous probability distributions that is defined on the interval [0,1]. It is characterised by two positive shape parameters α and β which are introduced as exponents of the random variable and control the shape of the distribution. An attractive feature of this distribution is the relationship that exists between the mean and the variance. While a normal distribution can have any variance, in the case of Beta distribution the moments are a function of the proportion (to be explained later). Consequently, the regressions that involve data from the unit interval (rates and proportions) tend to be heteroscedastic. In other words, the variance is smaller at the extremes 0 and 1, while it becomes larger when we approach 0.5. The moments (mean and variance) of this distribution are:

$$E_x = \frac{\alpha}{\alpha + \beta}$$

$$Var(x) = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$$

This means that we need to define suitable α and β that will have the following properties:

$$E(\frac{x}{m}) = \frac{x^*}{m}$$

$$Var(\frac{x}{m}) = \frac{\frac{x^*}{m}(1-\frac{x^*}{m})}{s}$$

where x is the actual and x^* the optimal allocation. The above properties are satisfied when For this solution it suffices to set:

$$\alpha = \frac{x^*}{m}(s-1)$$

$$\beta = (1-\frac{x^*}{m})(s-1)$$

with s being the *precision* parameter of the distribution that is to be estimated, which must be strictly greater than 1. The problem with these parameters is when the optimal allocation is a boundary portfolio (where all the optimal allocation is $x^* = 0$ or $x^* = m$) the variance is zero, fact that can create problems to the maximum log-likelihood[44] A proposed solution for this, is to define the parameters of the distribution in such a way that the variance never becomes zero. A pair of such parameters is:

$$\alpha = \frac{(bm/2 + (1-b)x^*)(s-1)}{m}$$

$$\beta = \frac{(m - bm/2 - (1-b)x^*)(s-1)}{m}$$

where b is the *bias* parameter and s is the *precision* of the distribution as before. Using these parameters, the moments we obtain are:

$$E(\frac{x}{m}) = b/2 + (1-b)x^*/m$$

$$Var(\frac{x}{m}) = -\frac{(2m - bm - 2x^* + 2x^*b)(-bm - 2x^* + 2x^*b)}{4m^2 s}$$

When $b = 0$ we obtain an unbiased estimator. In the following Figure the variance is plotted against the various possible values that the optimal allocation may obtain. It is apparent that the variance is not constant for all values of the allocation. It tends to zero at the extremes (0 and m) and is maximised at the point m/2.

The probability density function (*pdf*) of the *Beta*-distribution is given by

$$f(x) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1}$$

where $\Gamma(.)$ the *Gamma* function: $\Gamma(z) = \int_0^\infty t^{z-1}\exp(-t)dt$. α, β represent the shape parameters for location and dispersion. More specifically, as Smithson et al. 2006, point out, α pulls the density toward 0 while β pulls it toward 1. The next step is to define the respective α and β parameters in order to write the likelihood function. For the first stage we make the following assumptions. For colour i, the proportion of the amount allocated to that colour over the available experimental income $\frac{x_i}{m}$ follows the Beta distribution. Then, for colour j, we assume that $\frac{x_j}{m-x_i}$ follows a Beta distribution. Writing it in this way, guarantees that x_i is bounded in the space [0,m] and x_j in the space [0,m-x_i]. Then, $\frac{x_k}{m}$ is just the residual $1 - \frac{x_i}{m} - \frac{x_j}{m-x_i}$. We can now write the formulas that provide the respective shape parameters. In the first stage, we assume that x_i/m follows the Beta distribution and that it is centered on the optimal allocation, thus x_i^*/m. Similarly, we assume that $x_j/(m-x_i)$ is Beta

[44] We present this analytically in chapter 3.

distributed centered on $x_j^*/(m - x_i^*)$. We remind the reader that x stands for the actual allocation and x^* for the optimal. The parameters for the distribution are:

- $\alpha_{x_i}(x_i^*/m), \beta_{x_i}(x_i^*/m)$ for x_i

$$\begin{aligned} \alpha_{x_i} &= \left(\frac{x_i^*}{m}\right)(s_i - 1) \\ \beta_{x_i} &= \left(1 - \frac{x_i^*}{m}\right)(s_i - 1) \end{aligned} \qquad (2.25)$$

- $\alpha_{x_j}(x_j^*/(m - x_i^*)), \beta_{x_j}(x_j^*/(m - x_i^*))$ for x_j

$$\begin{aligned} \alpha_{x_j} &= \left(\frac{x_j^*}{m - x_i^*}\right)(s_j - 1) \\ \beta_{x_j} &= \left(1 - \frac{x_j^*}{m - x_i^*}\right)(s_j - 1) \end{aligned} \qquad (2.26)$$

- x_k is the residual of $m - x_i - x_j$, where x_i, x_j are the amounts allocated in the first round.

In the second stage, the allocation happens between the two remaining colours therefore, it suffices to assume that one of the conditional allocations is distributed according to the Beta distribution and the other is the residual. Take for instance the case where the information "the ball is not i" is revealed. The conditional allocation is made between $x_{j\neg i}$ and $x_{k\neg i}$ and the available income is $m_{\neg i} = x_j + x_k = m - x_i$. Then, we assume that $x_{j\neg i}/m_{\neg i}$ is distributed according to the Beta distribution. The shape parameters are:

$$\begin{aligned} \alpha_{x_{j\neg i}} &= \left(\frac{x_{j\neg i}^*}{m_{\neg i}}\right)(s_{j\neg i} - 1) \\ \beta_{x_{j\neg i}} &= \left(\frac{x_{j\neg i}^*}{m_{\neg i}}\right)(s_{j\neg i} - 1) \end{aligned} \qquad (2.27)$$

This ensures that $x_{j\neg i}$ is bounded on the interval [0,$m_{\neg i}$]. Notice that in forming the shape parameters, it is assumed that the *precision* parameter has different indicator for each of the random variables ($s_{j\neg i}, s_{i\neg j}, s_{i\neg k}$) implying that it takes different values. As there is no story to explain behaviourally how this parameter is formed and in order to avoid overfitting by adding additional parameters, we assume that this precision parameter is the same at all stages and is equal to s. The rest of the conditional states follow the same specification.

2.4.3 Forming the Likelihood Function

In order to obtain the values of the parameters that best fit the model we need to specify the likelihood function that will be maximised. When the variables are rounded, then it is better to make the evaluation of a realising value, close to the measured value. This means, that for a random variable with value y, the probability that this variable takes values in the interval $[y-\varepsilon, y+\varepsilon]$ for ε small enough, is considered when forming the likelihood function. In our analysis, we consider this interval to be the $[\frac{x-0.5}{m}, \frac{x+0.5}{m})$ because the values were represented to subjects in integers. The latter allows us to use instead of the *pdf*, the *cumulative distribution function* (cdf). We need to take into consideration that when using the Cumulative Distribution Function for a discrete value, in fact we calculate an area, contrary to the use of the Density Function, where we obtain the point probability of the discrete value. Now the probabilities can be written as:

- Prob($x=0$)=$cdf[\frac{0.5}{m}]m$
- Prob($x=i$)=($cdf[\frac{i+0.5}{m}]$-$cdf[\frac{i-0.5}{m}]$)m
- Prob($x=m$)=1-$cdf[\frac{m-0.5}{m}]m$

and then the contributions to the likelihood can be written as:

- if $x=0$, $\ln\left(cdfBeta\left(\frac{0.5}{m}, \alpha, \beta\right)m\right)$
- if $0<x<m$, $\ln\left(cdfBeta((\frac{x+0.5}{m}, \alpha, \beta) - cdfBeta(\frac{x-0.5}{m}, \alpha, \beta))m\right)$
- if $x=m$, $\ln\left(1-cdfBeta\left(\frac{m-0.5}{m}, \alpha, \beta\right)m\right)$

We know that each allocation is bounded between zero and m and are integers. In that point we are specifying parameters for the simulation (which are to be estimated later) b, s for the beta distribution. Then we are able to write down the log-likelihood equation.

In the case where the subjects fell short of time, the software was programmed to allocate zero at each colour. Consequently, in this case this question is not included to the analysis and the contribution to the likelihood function is equal to zero. This did not happen too often (around 5 subjects, 1-2 problems per subject).

Now we present the general form of the likelihood function. Writing down the density function, we have:

$$f(y|\alpha,\beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} y^{\alpha-1}(1-y)^{\beta-1}$$

where $y \in [0,1], \alpha, \beta > 0$ and $\Gamma(.)$ the gamma function. In the model to be estimated there are two random variables that are distributed according to the beta distribution. x_1, x_2 and x_3 as the residual. Assuming ϑ_0 a vector of parameters to be a estimate, the joint density function is:

$$f(x|\vartheta_0) = \prod_{i=1}^{n} f_i(x_i;\vartheta_0)$$

The likelihood function for the distribution is:

$$\mathcal{L}(\vartheta_0, x) = \prod_{i=1}^{n} \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1}$$

$$= \left(\frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}\right)^{n} \prod_{i=1}^{n} x_i^{\alpha-1} \prod_{i=1}^{n} (1-x_i)^{\beta-1}$$

and for simplicity and computation easiness we can write the log-likelihood as:

$$\ln \mathcal{L} = n\ln(\Gamma(\alpha+\beta)) - n\ln(\Gamma(\alpha)) - n\ln(\Gamma(\beta)) + \tag{2.28}$$

$$(\alpha-1)\sum_{i=1}^{n}\ln(x_i) + (\beta-1)\sum_{i=1}^{n}\ln((1-x_i)) \tag{2.29}$$

where n the number of observations. By maximising the maximum likelihood function $\mathcal{L}(\vartheta_0, x)$ we choose from the parameter vector ϑ_0, a ϑ that *maximises the likelihood* of observing the actual sample. Assuming that the observations are *independent* and *identically* distributed, the probability of obtaining the set of observations that the data consists of given the specified model is equal to the product of the probabilities of obtaining each value.

2.4.4 Obtaining the Standard Errors

For the econometric analysis of the data the general non-linear optimisation package[45] *rsolnp* developed by Ghalanos and Theussl (2012) was used. This solver implements a general nonlinear *augmented Lagrange multiplier* method. More specifically, in the analysis, it was applied using the *gosolnp* option. This option allows for *random initialisation* of the starting parameters as well as multiple restarts of the solver so as to ensure that the solution is a global optimum. As the problem is quite complex in nature, it is expected that the likelihood surface will not be smooth and consequently, a global maximum will not be easy to reach. If one is constrained to use a specific set of starting parameters, there is no guarantee that the searching algorithm will not stop at a local maximum and return this as the optimal point. Using random starting values for the parameters and multiple starting points, the chance of obtaining a local maximum is significantly reduced and as the number of random values as well as multiple starts is increased, this chance can be eliminated (always under the respective computing process time cost). As these methods are gradient-free, it is impossible to compute standard-errors directly and one needs to use computational methods in order to do so. The standard way requires the inversion of the covariance matrix and then the standard errors are given by the square root of the diagonal. Nevertheless, this method is not always possible to work, as there are singularity issues that make it impossible to inverse the matrix leading to unreliable results. For this reason, in the next section we do not present the standard errors.

2.4.5 Number of Parameters

Finally, we present the number of parameters for each model, as well as the constraints that we impose for the upper and lower bounds of the estimators. The goodness of fit of the various models depends on the number of parameters that need to be estimated. If this number high, the models is suffering from *overfitting*. Overfitting appears when a statistical model describes random error

[45]Package for the *R* programming language for statistical computing[46]

or noise instead of the underlying relationship of the training set. Effort has been made to take the latter into consideration. On the one hand, the number of the parameters to be estimated was held to minimum. On the other hand, when reducing the number of parameters was not possible, the appropriate criteria have been applied to allow comparison among different models (comparing the models using the Akaike and the Bayesian Information Criterion) as is described later. Table 2.12 presents the number of parameters that were estimated for each of the models.

Table 2.12 Models and Number of Parameters to Estimate

	Preferences	**Type**	**Updating Rule**	**# Params**	**Parameters**
1	SEU	-	Bayesian	4	p_i, p_j, r, s
2	MEU	Resolute	-	5	$\underline{p_i}, \underline{p_j}, \underline{p_k}, r, s$
3	MEU	Naive	GBU /MLE	5	$\underline{p_i}, \underline{p_j}, \underline{p_k}, r, s$
4	MEU	Sophisticated	GBU/MLE	5	$\underline{p_i}, \underline{p_j}, \underline{p_k}, r, s$
5	CEU	Resolute	-	8	$v_i, v_j, v_k, v_{ij}, v_{jk}, v_{ik}, r, s$
6	CEU	Naive	OPT	8	$v_i, v_j, v_k, v_{ij}, v_{jk}, v_{ik}, r, s$
7	CEU	Naive	GBU	8	$v_i, v_j, v_k, v_{ij}, v_{jk}, v_{ik}, r, s$
8	CEU	Naive	DS	8	$v_i, v_j, v_k, v_{ij}, v_{jk}, v_{ik}, r, s$
9	CEU	Sophisticated	Optimistic	8	$v_i, v_j, v_k, v_{ij}, v_{jk}, v_{ik}, r, s$
10	CEU	Sophisticated	GBU	8	$v_i, v_j, v_k, v_{ij}, v_{jk}, v_{ik}, r, s$
11	CEU	Sophisticated	DS	8	$v_i, v_j, v_k, v_{ij}, v_{jk}, v_{ik}, r, s$
12	SCEU	Resolute	-	5	p_i, p_j, α, r, s
13	SCEU	Naive	Bayesian	5	p_i, p_j, α, r, s
14	SCEU	Sophisticated	Bayesian	5	p_i, p_j, α, r, s

The minimum number of parameters is required by the SEU model (4), while the maximum by the CEU (8). For all the parameters that express *beliefs* (probabilities, capacities or lower bounds of probabilities) the upper limit was set to 1 and the lower to 0. For the SEU model, we require that all the probabilities sum up to 1. For the MEU this constraint does not hold any more. If it does, then there is no distinction between the two models, SEU and MEU. For the CEU model, we require that the individual capacities are no larger than the joint capacities. That is to say, if v_i, v_j are the capacities for states i and j and v_{ij} is the joint capacity, it should hold that $v_{ij} - v_i \geq 0$ and $v_{ij} - v_j \geq 0$. The same holds for the rest of the combinations, for all the possible states of the world. The parameter of risk aversion r, is having the same bounds for all the specifications we consider. As the utility function is of the CRRA form, the coefficient can take both positive and negative values. When $r > 1$, then the decision maker is risk loving. Our design is not appropriate to define this type, thus we constraint the parameter to the interval [0,1]. In addition, we do not consider negative values for r, as this indicates very high levels of risk aversion. The precision parameter s is assumed to take values between 1 and 60. The largest the precision, the less noisy the dataset is. Finally, for the case of SCEU, the weighting parameter α takes values in the interval $(0,1]$ with 1 indicating a SEU decision maker.

2.5 Results

Before presenting the results, it is important to stress the fact that all the analysis, directly depends on the assumptions that we made in the previous section. More specifically, we make assumptions regarding the form of the preferences of the decision makers, the representation of their utility function as well as the stochastic part of their decision process and objective is to obtain some insight on which of the above combinations provides a better explanation of how people behave on a sequential problem under ambiguity. In addition, an assumption that is implicitly made is that the type of the subjects remains constant during the experimental session and the same holds for their preferences. The analysis was conducted at the individual level. For each of the 58 participants, we fitted all the possible combinations that we described in section 2.2.[47] For each subject, for each preference functional and for each specification, we have estimates of the parameters of the functional (probabilities, lower bounds of probabilities or capacities,weighting parameter), of the risk aversion parameter (r), the precision parameter s and the value of the *maximised log-likelihood.*

We start presenting the results with some descriptive statistics. As the main question is whether subjects comply to either Dynamic Consistency or Consequentialism, we first use the *strict* definition of *dynamic consistency* that was adopted in section 2.2, which requires that the allocations at stage 1 and stage 2 should be exactly the same (the proof of this proposition is provided in section A.1 in Appendix A). In Table 2.13 presents the percentage of subjects whose choices in the

[47]The subject-level analysis created a large data-set that contains the estimated preferences for each individual. We do not report it here, but full details are available upon request.

second stage were different than what they have allocated to the first stage.

Table 2.13 Percentage of Changes

	Not Blue	Not Yellow	Not Pink
$N = 58$	70.54%	72.98%	71.11%
std	(0.2)	(0.21)	(0.25)
Min (subject 58)	50%	0.09%	0%
Max (subject 23)	100%	94%	94%

On average, 70.54% of the subjects (42 out of 58) had allocations that were different at each stage. If one expects that no change should be done on the initial allocations, then this percentage shows an extremely high violation of the axiom of dynamic consistency. Nevertheless, this measure is quite sensitive as even if a subject is making the same allocations at both stages and has a different allocation in only one problem, then this measure categorises this subject as dynamically inconsistent. This is quite extreme, since as it can be seen in the third row of Table 2.13, subject 58 was dynamically consistent in the cases where the ball was not Yellow or Pink with almost zero percentage of changes. Interestingly though, 50% of the times where the information was that the ball is not Blue, this subject was changing her or his initial allocations (recall that Blue was the colour with the lowest probability). On the other hand, subject 23, was the subject with the maximum number of changes in the allocations, where almost at all problems she or he changed the initial allocations. Figure 2.7 illustrates the percentages of changes with the respective bar-charts.

As the previous measure is quite extreme and not very informative, it is interesting to obtain an illustration of how the decisions look like. In order to obtain an idea of the different behavioural patterns and the degree of *heterogeneity* that we observed in the lab, it is helpful to focus on few representative scatter-plots. All scatter-plots show the decisions of a subject for both states and for each conditional state. Figures 2.8-2.12 show the decisions of five subjects for the three conditional states of the world. Let us focus on the right hand side graph, where the conditional state $\neg Pink$ is illustrated. On the vertical axis, the payoff when the ball is Blue is represented while on the horizontal the respective payoff when the ball is Yellow. The *red* dots stand for the allocation at period 1, while the *blue* stars

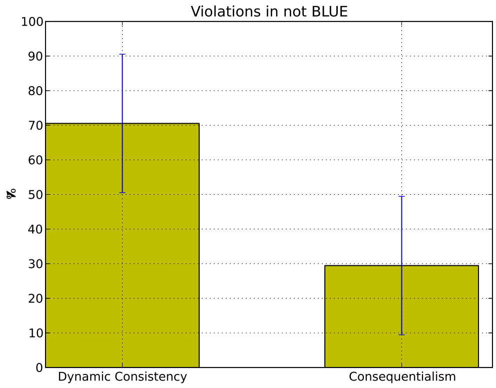

(a) Ball is Not Blue

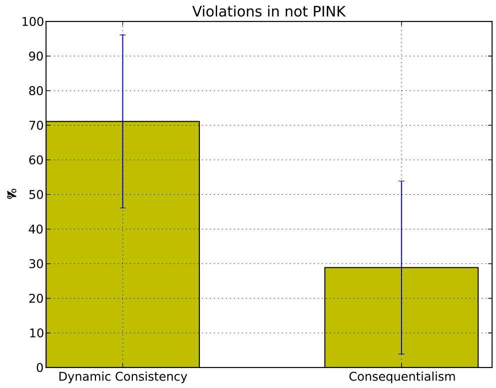

(b) Ball is Not Pink

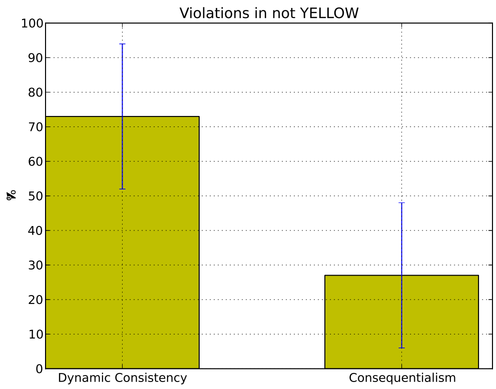

(c) Ball is Not Yellow

Figure 2.7 Percentage of Changes in the Three Conditional States

for the conditional allocation[48] The 45 degrees line, represents all the allocations where the decision maker secures a sure payment no matter the real state of the world, behaviour that shows extreme risk aversion. Figure 2.8 is exactly the behaviour that ones expects from a Resolute decision maker that somehow commits to the same decision at both stages.

Other patterns that have been observed, include Figure 2.11 where the subject initially allocates nothing or almost nothing to an asset and upon the receipt of the partial information, the allocation becomes more ambiguous averse, and gets closer to the equal allocation line. A different motive that worths mentioning, is the one at Figures 2.9 and 2.10 where after the receipt of information, the subjects

[48]Roughly 20 observations for each conditional state.

show some kind of ambiguity aversion and tend to make allocations that are close to the equal allocation line. 2.11 presents similar behaviour but in a more extreme way. Finally, there are a few subjects who show extensively high degrees of preference reversals. Figure 2.12 shows the case where all the income is allocated to one asset and at the second stage, the allocation to this asset is zero and all the income is allocated to the remaining asset. The majority of the subjects realised allocation that were similar to the pattern that is illustrated as Figures 2.9 ad 2.10.[49].

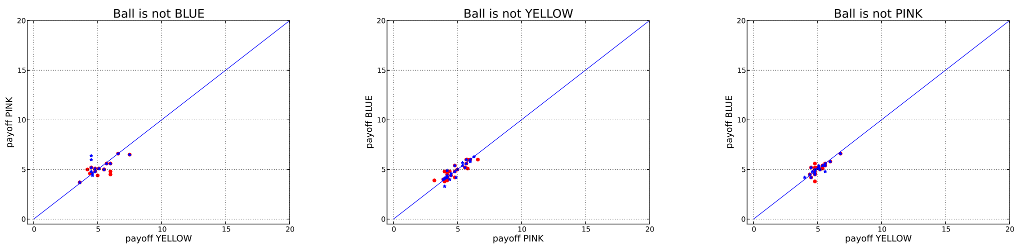

Figure 2.8 Subject 4

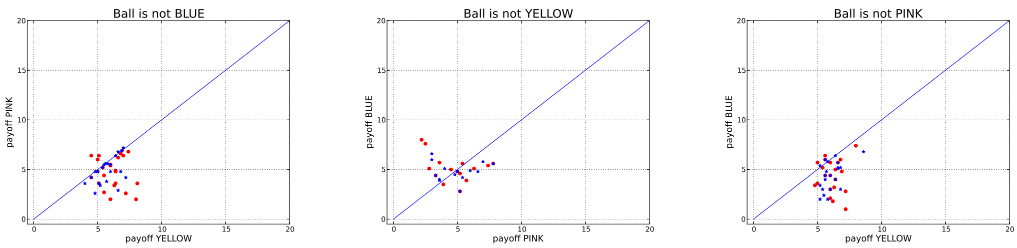

Figure 2.9 Subject 22

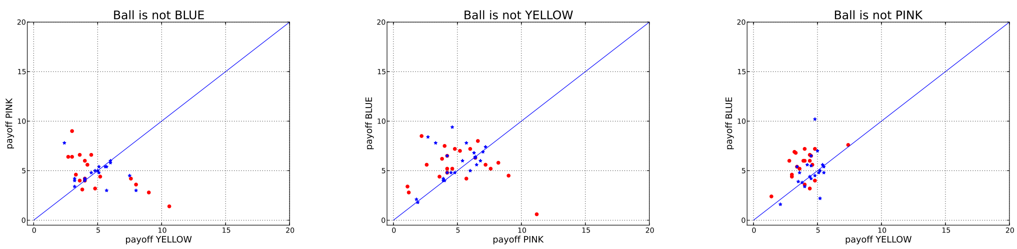

Figure 2.10 Subject 27

Table 2.14 shows the average values of the estimated parameters for the 14 specifications. Recall that the number of parameters differs for each specification and the only common parameters for all the different preference functionals, are

[49]The patterns that presented above have been repeated by many participants and this is the reason why we do not include all the plots as an appendix. The full set of scatter-plots is available upon request.

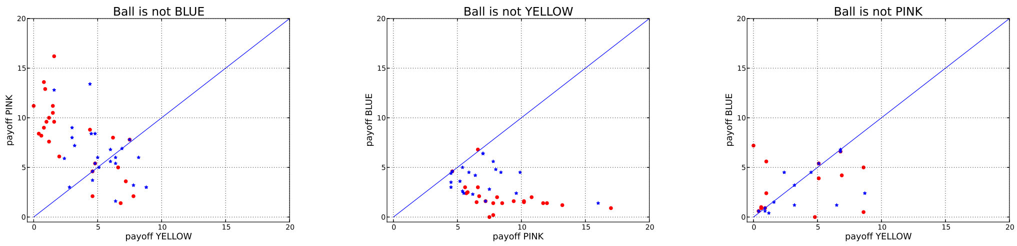

Figure 2.11 Subject 45

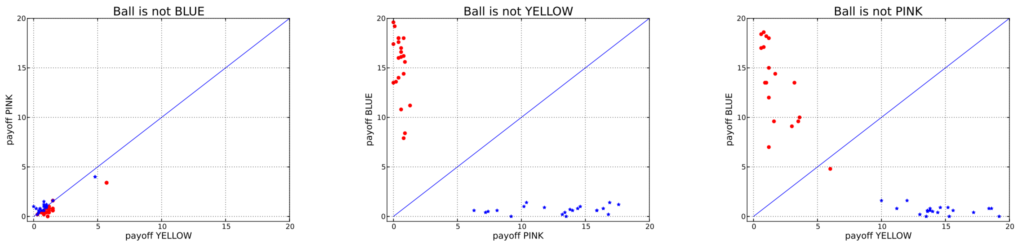

Figure 2.12 Subject 48

the coefficient of risk aversion r and the precision parameter s. The estimated value of the risk parameter is on average equal to .24 for all the specifications. This is in line with the usual findings, where subjects are found to be risk averse. The precision parameter s is an indicator of the noise in the data. The largest the value of the parameter, the less noisy the data are. In the experiment, an average value of around 24.5 was estimated which is significantly different than one. Then, the rest of the parameters represent beliefs. In the case of SEU, first row in Table 2.14, the beliefs are additive and correspond to the subjective beliefs of the decision maker. Recall that i,j,k stand for the three colours that represent assets. i is for Blue, j is for Yellow and k is for Pink. Also recall that the actual probabilities during the experiment were .2 for Blue, .5 for Yellow and .3 for Pink. Notice, that although the estimated parameters preserve the correct magnitude order for the probabilities (Yellow is larger, then Pink, then Blue) they are far from their actual values. More specifically, the probabilities for the less likely events (Blue and Pink) are over-weighted, while the probability for the state of the world that is most likely to happen, is undervalued. This can be partly explained by the aversion towards ambiguity that characterises subjects' beliefs.

However, this Table does not provide much information on how these specifications can be compared to each other, regarding the *goodness of fit*, due to the

Table 2.14 Estimates of the Parameters

Specification	p_i	p_j	p_k	r	s	α	LL
SEU	0.288	0.362	0.350	0.210	23.93	-	-487.86
MEU_{res}	0.231	0.341	0.303	0.281	24.34	-	-459.526
MEU_{naive}	0.250	0.334	0.307	0.277	24.50	-	-456.171
MEU_{soph}	0.257	0.333	0.265	0.253	21.44	-	-472.250
$SCEU_{res}$	0.280	0.372	0.348	0.208	23.87	0.852	-460.640
$SCEU_{naive}$	0.288	0.361	0.350	0.228	24.30	0.815	-458.307
$SCEU_{soph}$	0.288	0.362	0.350	0.201	24.09	0.983	-461.306

Specification	v_i	v_j	v_k	v_{ij}	v_{jk}	v_{ik}	r	s	LL
CEU_{res}	0.245	0.298	0.177	0.597	0.681	0.526	0.254	25.03	-454.03
$CEU_{naive}OPT$	0.211	0.316	0.241	0.595	0.662	0.550	0.260	24.88	-449.25
$CEU_{naive}DS$	0.249	0.318	0.212	0.594	0.679	0.510	0.248	24.93	-449.49
$CEU_{naive}GB$	0.253	0.334	0.260	0.615	0.672	0.541	0.244	24.66	-451.23
$CEU_{soph}OPT$	0.218	0.325	0.242	0.601	0.662	0.548	0.251	24.89	-456.22
$CEU_{soph}DS$	0.232	0.312	0.239	0.595	0.645	0.546	0.243	24.88	-454.94
$CEU_{soph}GB$	0.215	0.320	0.241	0.598	0.662	0.549	0.256	24.88	-454.735

different number of *degrees of freedom*. We proceed by correcting the value of the log-likelihoods, taking into consideration both the number of the parameters and the size of our sample.

2.5.1 Bayesian and Akaike Information Criteria

The models are compared on the basis of the value of the *maximised log-likelihood*. Nevertheless, a standard criticism against the non-Bayesian models is that it is natural to perform better (when they do) due to the additional parameters, due to *overfitting*. As was discussed before, we estimate different preference functionals, the parameters' number of which, was not always the same. For this reason, also compare the models using the *Akaike* Information Criterion (AIC) and the *Bayesian* Information Criterion (BIC). Both criteria are applied for model selection and they penalise the models according to the number of the free parameters that they allow. In both criteria, a *lower* value indicates a better fitting. More analytically, the formulas for the two are provided below:

$$BIC = -2\ln(L(\hat{\theta}|x)) + k\ln(n)$$

$$AIC = -2\ln(L(\hat{\theta}|x)) + 2k$$

where $ln(L(\hat{\theta}|x))$ is the value of the maximised log-likelihood, k is the number of the free parameters in the model and n the number of observations. In the models that we are testing, the value of the likelihood is different for each of the subjects, the number of the parameters changes depending on which model we focus on (for SEU there are 4 parameters, for MEU 5, for SCEU 5 and for CEU 8) and the number of observations is constant and equal to 120 (60 + 60) for all subjects.

In addition, we also use the *corrected version* of the AIC (AIC_C) for finite samples, when the size of the sample n is small, or when the number of the parameters k is large enough. A rule of thumb, suggests that the AIC_C should be used when the ratio n/k is less than 40. The formula for the AIC_C is:

$$AIC_C = AIC + \frac{2k(k+1)}{n-k-1}$$

Table 3.15 presents the respective values for BIC, AIC and AIC_c.

Table 2.15 Corrected Log-Likelihoods

	SEU	MEU_{res}	MEU_{naive}	MEU_{soph}	$SCEU_{res}$	$SCEU_{naive}$	$SCEU_{soph}$	Obs
Uncorrected LL	-487.86	-459.52631	-456.170862	-472.25	-460.64	-458.307	-461.3	58
	(95.99)	(100.88)	(99.89)	(94.52)	(98.71)	(96.33)	(95.53)	
BIC	944.28	942.99	936.28	968.44	945.22	940.55	946.54	58
	(191.98)	(201.77)	(199.79)	(189.05)	(197.43)	(192.66)	(191.06)	
AIC	933.13	929.05	922.34	954.5	931.28	926.61	932.61	58
	(191.98)	(201.77)	(199.79)	(189.05)	(197.43)	(192.66)	(191.06)	
AIC_C	933.48	929.58	922.87	955.03	931.81	927.14	933.13	58
	(191.98)	(201.77)	(199.79)	(189.05)	(197.43)	(192.66)	191.06	

	CEU_{res}	$CEU_{naive}OPT$	$CEU_{naive}DS$	$CEU_{naive}GB$	$CEU_{soph}OPT$	$CEU_{soph}DS$	$CEU_{soph}GB$	Obs
Uncorrected LL	-454.034	-449.246241	-449.492103	-451.226	-456.22	-454.94	-454.735	58
	(101.37)	(100.88)	(97.99)	(100.36)	(98.16)	(98.25)	(100.92)	
BIC	946.37	936.79	937.28	940.75	948.23	946.95	948.08	58
	(202.76)	(201.78)	(195.99)	(200.32)	(201.42)	(197.04)	198.05	
AIC	924.07	914.49	914.98	918.45	926.73	926.88	928.40	58
	(202.76)	(201.78)	(195.99)	(200.72)	(196)	(198.93)	(197.19)	
AIC_C	925.37	915.79	916.28	919.75	926.20	928.32	926.10	58
	(202.76)	(201.78)	(195.99)	(199)	(195.6)	(195.96)	201.64	

The Table contains the values for the Uncorrected Log-Likelihoods and the Corrected Log-Likelihoods (BIC, AIC, AIC_c)

We are presenting the results model by model and then we generalise the discussion to the total number of possible combinations. Assuming that subjects adhere to only one non-EU preference functional (e.g. MEU), we test within this preference functional, how well the different types we defined fit the data.

2.5.2 MaxMin Expected Utility (MEU)

We start with the MaxMin model. In this specification, we assumed three different types of decision makers. Recall that in this specific preference functional, and with the specific choice task, it is not possible to distinguish between the different updating rules. As the objective is to test how well non-EU models fit the data, it is reasonable to compare this model using SEU as a benchmark. We compare the three models using the corrected log-likelihood and more specifically the BIC criterion. Using the Bayesian Information Criterion, we rank the three specifications from the best fitting to the worst. Based on the rankings for each of the subjects, Table 2.16 shows the cumulative percentages of subjects where the specifications are ranked as best, second best and so on.

Table 2.16 MEU Rankings Based on the BIC

Specification	1	1-2	1-3	1-4
SEU	2	9	33	100
MEU_{res}	26	78	97	100
MEU_{naive}	66	94	100	100
MEU_{soph}	7	21	71	100

All values represent cumulative percentages

It is interesting that MEU_{naive}, is ranked first for 66% of the subjects and for 94% is ranked first or second. $MEU_{resolute}$ is ranked first more than 1/4 of the total population of the experiment. Finally, for less than 10% of the subjects, SEU is ranked first or second.

An indirect test that we can apply at this stage in order to compare the SEU and the MEU specifications is to check whether the underlying probabilities satisfy the additive property of the SEU. In Figure 2.13, histograms of the sum of probabilities for each of the different types are illustrated. In the case of SEU this sum is equal to 1. From the Figure it is obvious that for the resolute and the sophisticated type there is heterogeneity regarding the beliefs, while in the case of the naive decision maker, it seems that for almost 50% of the subjects, the sum of the probabilities is very close to one, but only for 16 subjects (28%) it is exactly equal to 1. Regarding the magnitude of the probabilities, they retain the correct order (Yellow highest, and so on) but they are significantly lower that in the SEU case, as they represent the lower bounds of the probabilities.

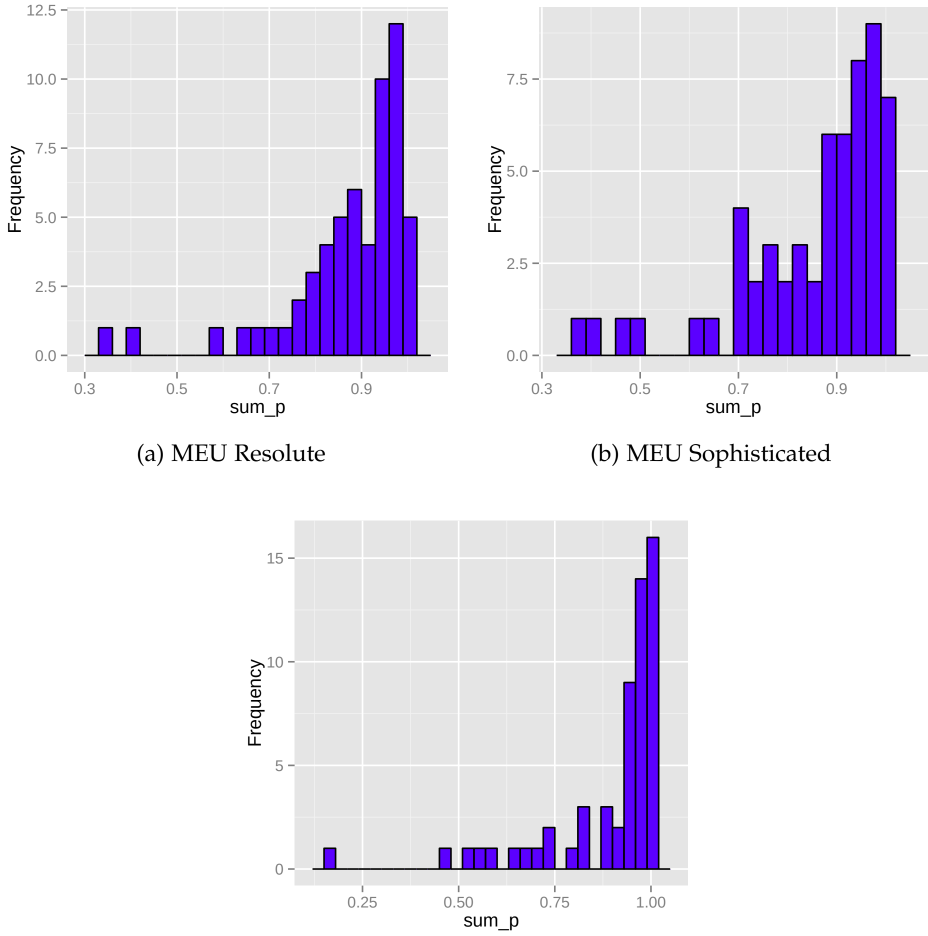

(a) MEU Resolute

(b) MEU Sophisticated

(c) MEU Naive

Figure 2.13 Sum of the Lower Bound Probabilities for MEU

2.5.3 Source Choquet Expected Utility (SCEU)

Table 2.17 shows the ranking for the different types of the $SCEU$ model, including SEU, according to the BIC criterion. $SCEU_{soph}$ is ranked first for 47% of the subjects and either first or second for 81%. Similarly, $SCEU_{naive}$ is ranked first or second for 75% while SEU only for 2% of the subjects.

Table 2.17 SCEU Rankings Based on the BIC

Specification	1	1-2	1-3	1-4
SEU	0	2	2	100
$SCEU_{res}$	19	41	93	95
$SCEU_{naive}$	34	75	100	100
$SCEU_{soph}$	47	81	109	109

All values represent cumulative percentages

In Figure 2.14, the histograms of the parameter α for the different types are illustrated. Recall that α is a weight parameter and when its value is equal to 1, SCEU collapses to SEU as there is no weighting of the probabilities.

In the case of a resolute decision maker, the α parameter is quite dispersed. The same holds for the naive type to a lesser extent. On the contrary, for more than 50% of the subjects, the value of this parameter tends to 1, where the model collapses to SEU. Regarding the magnitude of the parameters' values that represent beliefs, are quite close to the values of the SEU.

2.5.4 Choquet Expected Utility (CEU)

Tables 2.18, 2.19 and 2.20 show the same kind of classification for the case of CEU and for the the three updating rules, optimistic, Dempster-Shafer and Generalised Bayesian rule respectively.

Table 2.18 CEU_{OPT} Rankings Based on the BIC

Specification	1	1-2	1-3	1-4
SEU	7	10	15	100
CEU_{res}	31	52	92	100
CEU_{naive}	41	57	100	100
CEU_{soph}	21	81	98	100

All values represent cumulative percentages

CEU_{naive} is classified first for both the optimistic updating rule and the Dempster-Shafer (41 and 43%). On the contrary, when the Generalised Bayesian rule is applied, the best fitted specification is the resolute type with 33%.

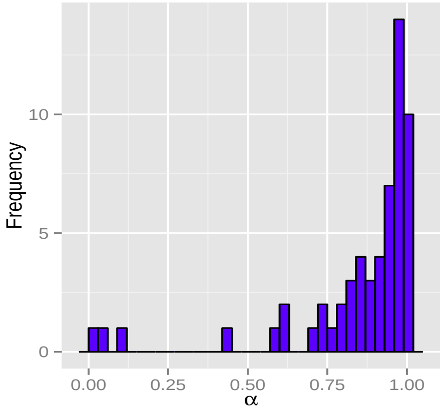

(a) SCEU Resolute

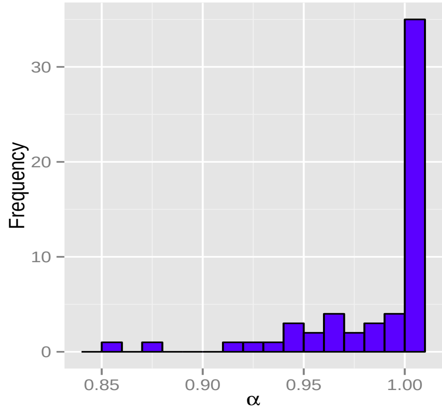

(b) SCEU Sophisticated

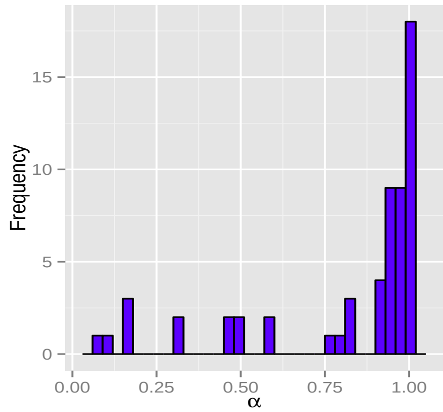

(c) SCEU Naive

Figure 2.14 α Parameter for SCEU

Table 2.19 CEU_{DS} Rankings Based on the BIC

Specification	1	1-2	1-3	1-4
SEU	7	12	17	100
CEU_{res}	29	43	91	100
CEU_{naive}	43	67	100	100
CEU_{soph}	21	78	100	100

All values represent cumulative percentages

Regarding the magnitude of the parameters' values, they are preserve the correct ranking, but they are far from the actual values. This is expected as they represent capacities that are not additive.

Finally, we aim to classify all 14 specifications for all the subjects depending on the BIC criterion. $SCEU_{soph}$ is ranked first for 25% of the subjects and first

Table 2.20 CEU_{GB} Rankings Based on the BIC

Specification	1	1-2	1-3	1-4
SEU	7	12	15	100
CEU_{res}	33	52	92	100
CEU_{naive}	29	57	100	100
CEU_{soph}	31	79	95	100

All values represent cumulative percentages

or second for 37%. $CEU_{naive}OPT$ is ranked first for 16% and $SCEU_{naive}$ first or second for 25%. The SEU fails to be classified first for any of the subjects while it is classified second for 5%. The results are significant for most of the subjects at both 1% and 5% significance levels, based on a *Likelihood Ratio*[50] test.

2.6 Conclusion

When the preferences are assumed to be Subjective Expected Utility, it is automatically assumed that two conditions are satisfied. On the one hand, it is expected that the decision maker is probabilistically sophisticated, that she is able to attach a subjective probability distribution to the various events of the world. On the other hand, it is assumed that these subjectively formed priors, are updated according to the Bayesian rule when additional information regarding the future states of the world becomes available. The combination of the two ensures that the decisions of the agent, simultaneously satisfy the axioms of Dynamic Consistency and Consequentialism. Nevertheless, based on the evidence that experiments inspired by the Ellsberg paradox provide, subjects usually do not satisfy both axioms at the same time. Consequently, SEU is not a suitable model anymore and the issue of how people behave in a dynamic problem under ambiguity, becomes an issue of considerable controversy. Yet, the literature has not converged to a common accepted answer, regarding which axioms are being satisfied or violated and how the prior beliefs are updated in the absence of an objective probability distribution.

[50] The Likelihood Ratio Test (LRT) is used to test for significance when two models are *nested* (one model is reduced to the other of a linear constraint is imposed to the parameters). In our case we compare every model with SEU, as an appropriate constraint either to the additive property of beliefs or to the weighting function for the probabilities, can transform each model to SEU.

In this chapter we present an economic experiment on dynamic decision making under ambiguity. By employing a Bingo Blower as the source of ambiguity, we ask subjects allocation questions in a sequential problem. The environment is sufficiently ambiguous and the experimental task is designed in such a way that it allows us, by making several structural assumptions, to fit different specifications regarding preference functionals and updating rules. The objective of the experiment is twofold. First, we aim to understand which axioms are being satisfied in an appropriate decision task. Then, accepting that agents have non-EU preferences, which may lead to dynamic inconsistencies, we try to understand how people decide and how they resolve these possible inconsistencies by defining three different types (resolute, naive and sophisticated). From the analysis, we find that SEU lacks descriptive power for this kind of problems. A considerable proportion of the experimental population seems to adopt the sophisticated type of decision maker, where the problem is solved backwards. Nevertheless, this strategy is characterised by ambiguity aversion, as the beliefs are weighted. An additional result, is that an different, almost same in size proportion, behaves in a naive way without realising the dynamic nature of the problem. It is natural to expect that these kind of behaviour has consequences to all aspects of economic life. Dynamic decision making includes savings decision, financial decisions and investment decisions, where the agents are asked to plan ahead in an environment that is characterised by incomplete information and the ambiguity regarding the future levels of income, rates of return and prices is always present. Thus, it makes one think, what is the cost of these inconsistencies on the welfare of the decision makers and what is the role of government intervention to 'nudge' and minimise this cost if it exists. Of course it is difficult to generalise and further research is needed. In the Conclusion we present the possible extensions of this experiment that will provide more evidence on how ambiguous beliefs are updated.

Chapter 3

Modelling Error Specification in a 3-Way Allocation Problem[1]

3.1 Introduction

Experimentalists are increasingly using allocation problems to make inferences about subjects' preferences - the reason being that allocation problems appear more informative than other types of problems - such as pairwise choices, Holt-Laury price lists Holt and Laury (2002), and the Becker-DeGroot-Marschak Becker et al. (1964) mechanism. At the same time some experimentalists are broadening the type of allocation problem, moving from allocations over just two events to allocations over more than two, again to get more information from experiments. Even with just two allocations, the issue of the error process is already interesting; going to allocations over more than two increases the interest as well as the complexity of the problem. In the previous chapter, we needed to made a simplification assumption, that the utility function of the subjects is represented by a CRRA utility function, in order to overcome the problems the boundary allocations create to the calculation of the log-likelihood function. We have seen that in the case where the optimal allocation is zero or negative, the specification that we applied degenerated and consequently an estimation of the underlying parameters is not feasible. Nevertheless, assuming a power function to represent utility, is

[1]This chapter is based on joint work with John Hey and Xueqi Dong. This research was funded by the Super Pump Priming Fund, Department of Economics, University of York (RIS 50) awarded to John Hey.

a quite restricting assumption as it rules out behavioural patterns that are reasonable to be expected during an experimental session or even in real-life economic applications. In this chapter, motivated by the latter, we present two alternatives to the CRRA modelling specifications for the stochastic term, that allow for boundary allocations. We run an extended simulation to obtain some insight of what happens when the true specification according to which the data are generated differs from the one assumed for the econometric estimation. Based on the results of the simulation, we present the design and the results of an economic experiment on decision making under risk, suitably designed to allow for comparison between the proposed stochastic specification. We find that when the wrong specification is assumed in order to analyse the data, there may appear considerable effects from this mis-specification. We also find, that when the degree of risk aversion is significantly high, the best specification is the one that preserves CRRA to represent utility.

The chapter is organised as follows: in the next section, we present the decision model upon which the chapter is based, in section 3.3 we present the three different specifications that we propose, in section 3.4 we present the results of an economic simulation we run to compare the combinations between true and estimated specifications. In section 3.5 the design and the results of the experiment are discussed. In section 3.6, the assumptions regarding the econometric analysis are illustrated. In the following section the results of the experiment are discussed. Then we conclude.

3.2 The Decision Problem

As the objective is to obtain data on allocation problems *per se*, it is necessary to have the subjects undergo a task as simple as possible. This is a twofold necessity. On the one hand it provides clean, robust data that is able to provide insights on the question under investigation. On the other hand, the simplicity of the task enables us to use simple methods of analysis that do not develop elaborated models of decision making with a very large number of parameters to estimate, as was the case in chapter 2. For this reason, the decision problem is quite similar to what has been used in the previous experiment, but free of the dynamic/sequential dimension as well as the ambiguous environment. We consider a 3-way allocation problem, in which subjects are asked, in a series of problems (one of which will be randomly selected at the end of the experiment to determine the subject's payment) to allocate a given amount of tokens (experimental income) between three risky events, with given exchange rates[2] between tokens and money for each state, and with given probabilities for each state. We define a state space $S = \{s_i, s_j, s_k\}$ for the three different states of the world. By m we denote the quantity of tokens that is endowed in order to allocate by e_i, e_j, e_k the exchange rate for the three states and by p_i, p_j, p_k the respective probabilities. Let x_i, x_j, x_k be the three allocations that should always satisfy the budget constraint $x_i + x_j + x_k = m$. Assuming Expected Utility preferences, the subject's decision is to choose the allocations to maximize $\sum_{s=i}^{S} p_s u(e_s x_s)$ subject to the budget constraint. The first order conditions are $p_s e_s u'(e_s x_s^*) = \lambda$ where λ stands for the Lagrangian multiplier and x_s^* denotes the optimal allocation. After making all the allocations, a randomly selected problem is played out for real and the subject is paid the money equivalent (given the exchange rates of that problem) of the number of tokens allocated to that state by that subject.

The experiment will provide observations of the allocations that the subjects actually made. These observations will be used to infer the preference functions of the subjects. Specifically, it might be the case that it is desired to infer whether these preferences are CARA (Constant Absolute Risk Aversion) or CRRA (Constant Relative risk Aversion). In addition, it is desired to infer the degree of risk aversion, which is captured by the parameter r of the utility function. The CRRA function is defined as follows:

$$u(x) = \begin{cases} x^r & \text{if } r > 0 \\ \ln(x) & \text{if } r = 0 \\ -x^r & \text{if } r < 0 \end{cases}$$

and respectively the CARA utility function:

[2] The exchange rates are used in the exactly way as in chapter 2.

$$u(x) = \begin{cases} -\exp^{-rx} & \text{if } r > 0 \\ x & \text{if } r = 0 \\ \exp^{-rx} & \text{if } r < 0 \end{cases}$$

where r is the risk aversion parameter and x is the payoff of the corresponding state. Table 3.1 presents the criteria that must be satisfied by the risk aversion parameter depending on the risk attitudes that characterise the decision maker for the CRRA function. Table 3.2 presents the same information for the CARA function. Figures 3.3-3.6 illustrate the shape of the utility function for different levels of risk aversion.

Table 3.1 Risk Attitudes for the CRRA Function

$u(x)$	$\frac{\partial u}{\partial x}$	$\frac{\partial^2 u}{\partial x^2}$	**Attitude**	**r**
x^r	rx^{r-1}	$(r-1)rx^{r-2}$	Averse	$r < 1$
$log(x)$	$1/x$	$-1/x^2$	Neutral	$r = 1$
$-x^r$	$-rx^{r-1}$	$-(r-1)rx^{r-2}$	Loving	$r > 1$

Table 3.2 Risk Attitudes for the CARA Function

$u(x)$	$\frac{\partial u}{\partial x}$	$\frac{\partial^2 u}{\partial x^2}$	**Attitude**	**r**
$-\exp^{-rx}$	$r\exp^{-rx}$	$-r^2\exp^{-rx}$	Averse	$r > 0$
x	1	0	Neutral	$r = 0$
$\exp^{-rx}$	$-r\exp^{-rx}$	$r^2\exp^{-rx}$	Loving	$r < 0$

For either preference functional and for any given value of the parameter r, it is possible to find the optimal unconstrained allocations (the analytical solutions are provided in Appendix B). It is in this part where the specificity of the CARA function becomes apparent which is also one of the main motivations behind this chapter. As was discussed in chapter 2 the CARA utility function, allows for zero and negative allocations to the assets. As this is a state that it is impossible to be implemented in the lab[3], it is reasonable to expect that subjects implement their optimal constrained allocations, which is actually what we observe by observing in the actual experimental data-sets. In order to capture this constrained maximisation decision making process, we have developed a case-specific algorithm

[3]It is impossible to make subjects lose money in a lab experiment. A possible solution to this, is to endow subjects with a fixed amount of money or tokens at the beginning of the experiment, so in any situation, the worst scenario would be to default. But again, this practice entails additional difficulties, as it is extremely difficult to define, if it can be defined, what this endowment level should be, because one can never provide a big enough "fixed amount of money".

(details on Appendix B).

3.3 Theoretical Framework

In this section we present the different stochastic specifications that we propose. Taking into consideration the noise that exists in the experimental data is advantageous in two ways. First, it is natural to make the assumption that in addition to the deterministic part of each theory, there is a stochastic component that cannot be explained by the former. In other words, when subjects make decisions, they do make mistakes and consequently, they do not implement their optimal strategy (either constrained on unconstrained depending on the case). Sometimes, it seems that the stochastic part of the decision process is more important than the assumed preference functional (Wilcox (2008)). Furthermore, when one wants to obtain estimates of the parameters of an underlying preference functional, it is inevitable not to make some assumptions on the stochastic part of the decision making process when particulars techniques are applied for the estimation. This is also the case of this chapter. The suggested methodology to estimate the respective parameters is by using *Maximum Likelihood Estimation* techniques which the way that they are applied, implicitly require that several assumptions (sometimes rather strong) regarding the distribution that characterises the random variables under investigation are satisfied.

As will be shown later, following the standard practice in the literature and assuming an additive error term that is normally distributed is not possible in our case due to technical difficulties. A significant part of the analysis is to understand the cognitive process that the subjects undergo when they make a decision (in the specific case an allocation under risk). Following the conventional methodology, it is assumed that the objective of an agent is to maximise some kind of preference functional (which is deterministically given by the underlying theories[4]). We assume the following *timing* regarding the way that the decisions are made. First, the subjects calculate their optimal, constrained allocations with error, so that the error is already added into the actual constrained allocations. The reason for this

[4]Expected Utility theory (EU) and Cumulative Prospect Theory (PT) are the two common specifications applied in decision making under risk. The assumption of the existence of a preference functional, excludes from the available stochastic theories the *Random Preferences* specification.

assumption is due to the fact that the majority of the experiments, the decision task is usually constrained to discrete allocations and as a consequence, it is not possible to detect whether this discretisation happens before or after errors have been added. In our case, we assume that subjects first arrive at continuous allocations with error and then discretise them. If the order of the events was the other way round, that subjects first make the discretisation to the optimal constrained allocation before the error is added, then different stories should have been applied.[5] The objective of the decision maker is to find the optimal allocation that maximises the Expected Utility, taking as given the probabilities and the exchange rates for the possible states of the world, subject to the budget constraint. The optimisation program can be written as:

$$\max_{x_s^*} \sum_{s=\{i,j,k\}}^{S} p_s u(e_s x_s) \tag{3.1}$$

$$\text{s.t.} \sum_{s=\{i,j,k\}}^{N} x_s = m \tag{3.2}$$

or applying it to the three-way allocation problem:

$$\max_{x_i^*, x_j^*, x_k^*} p_i u(e_i x_i) + p_j u(e_j x_j) + p_k u(e_k x_k) \tag{3.3}$$

$$\text{s.t. } x_i + x_j + x_k = m \tag{3.4}$$

3.3.1 Specification 1 - CRRA

One of the main assumptions is that these errors are built 'on top' of the optimal constrained allocations, rather than on the optimal unconstrained allocations. Given this, it is not possible to follow the conventional methodology, and by assuming normality, to add a normally distributed error term to the optimal allocations in order to get the actual allocations. The consequence of doing so, is that the implied actual allocations might violate the non-negativity constraints.

Since we focus on allocation problems with a fixed income m (which means that the allocations are constrained to the interval [0,m]), it is convenient to work

[5] Possible solutions include the Beta-distribution, the Beta with bias and the two-Betas, all of them are to be explained later.

with proportions allocated rather than integers. This means that the optimal allocations x_i^*/m (proportionally) will be inside the interval [0,1]. A natural error story that is appropriate to accommodate the above, is to assume that the random variables x_i^* follow the Beta distribution, centered on the optimal allocations x_i^*/m. This stochastic specification has been used before by Hey and Pace (2014), to model behaviour in a 2-way allocation problem but under ambiguity. As was also discussed in chapter 2, the Beta distribution requires the specification of the *shape parameters* which are denoted as α and β. A standard property of the shape parameters is that:

$$\text{Mean} = \frac{\alpha}{\alpha+\beta} \tag{3.5}$$

$$\text{Variance} = \frac{\alpha\beta}{(\alpha+\beta)^2} \tag{3.6}$$

The respective shape parameters are:

$$\alpha_i = \frac{x_i^*}{m}(s_i - 1) \tag{3.7}$$

$$\beta_i = (1 - \frac{x_i^*}{m})(s_i - 1) \tag{3.8}$$

it is guaranteed that:

$$E(\frac{x_i}{m}) = \frac{x_i^*}{m} \tag{3.9}$$

$$Var(\frac{x_i}{m}) = \frac{x_i^*(m - x_i^*)}{m^2 s_i} \tag{3.10}$$

where s_i is the *precision parameter* which is an indicator of the precision with which the subjects make their choices (the higher the more precise). In this specification there is no bias in the allocations and the expression for the variance implies that the spread of the distribution of the random variable is smaller the closer that the optimal allocation is to the bounds (0 or m) which seems to be a reasonable behavioural assumption implying that subjects make less errors towards the bounds (see Figure 3.1).

As we focus on allocations over three outcomes, it is important to illustrate how this can be modelled. Since there are three allocations to be done, it is convenient to focus only on the two and regard the third one as the residual. Assuming that there are three states of the world $s \in \{i,j,k\}$ we focus on x_i and x_j as $x_k = m - x_i - x_j$. If we assume that both $\frac{x_i}{m}$ and $\frac{x_j}{m}$ are distributed according to the Beta, there is no guarantee that the non-negativity constraint will be always satisfied[6]. In addition, there is no reason to assume that the allocation among the two different assets is totally independent. Thus, instead of ignoring the cases where $x_i + x_j \leq m$ there is an alternative way where we assume that x_i/m is distributed on the interval [0,m] with shape parameters $\alpha_i = \frac{x_i^*}{m}(s_i - 1), \beta_i = (1 - \frac{x_i^*}{m})(s_i - 1)$ with the respective mean and variance to be x_i^* and $\frac{(m-x_i^*)}{m^2 s_i}$. On the other hand, x_j/m is distributed on the interval $[0, m - x_i]$. This means that the shape parameters now are $\alpha_j = \frac{x_j^*}{m-x_i^*}(s_j - 1), \beta_j = (1 - \frac{x_j^*}{m-x_i^*})(s_j - 1)$. Similarly, the mean is $\frac{x_j^*}{m-x_i^*}$ of $\frac{x_j}{m-x_i^*}$ and the variance:

$$\frac{x_j^*(m - x_i^* - x_j^*)}{(m - x_i^*)^2(s_j - 1)^2 s_j}$$

Multiplying by $m - x_i^*$ we obtain the expressions for the mean and the variance for x_j:

$$E(x_j) = x_j^* \frac{m - x_i}{m - x_i^*} \tag{3.11}$$

$$Var(x_j) = \frac{x_j^*(m - x_i^* - x_j^*)(m - x_i^2)}{(m - x_i^*)^2 s_j} \tag{3.12}$$

The latter expression retains the properties of the variance discussed before. By 3.12 it seems that the mean of x_j is biased. Nevertheless, this depends on the value of x_i which when it is unbiased, it is guaranteed that the unconditional mean of x_j is equal to x_j^*. The method described above ensures two issues. Firstly, the non-negativity constraints are always satisfied. Then, the behavioural assumption that the allocation to one asset is not independent from that which is allocated to another asset in the same portfolio, is satisfied. If one wants to reinforce the results above, the assumption that the optimal allocations (the proportions) are always in the interval [0,1] must be satisfied at all times. Consequently, the utility function that is chosen to represent preferences plays a significant role. The simple result presented before, holds always when the utility function is bounded to the permitted values (in our case the interval [0,m]). This is true when the assumed func-

[6]This is crucial, as in the experiment not allocate negative amounts.

tion has the form of a Constant Relative Risk Aversion function (CRRA). Next we discuss some cases where the specification described above fails to capture some crucial behaviourial anomalies (mainly due to the assumption of non-negative and non-boundary allocations).

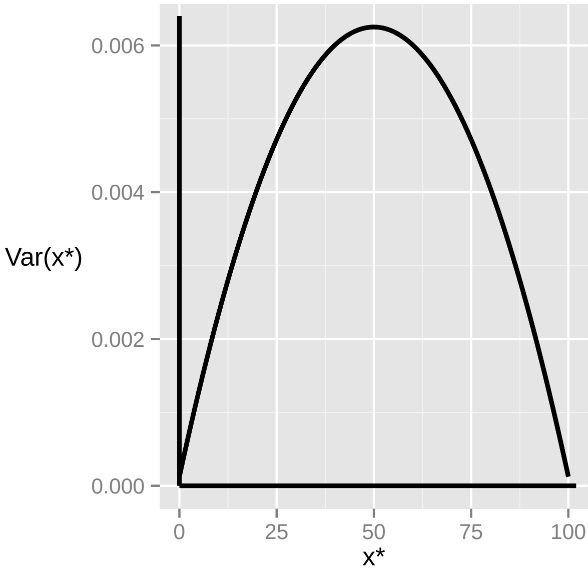

Figure 3.1 Variance of Beta Distribution with Bias, $\beta = .01$

The above specification is convenient when the underlying assumed utility function has a CRRA form. When the preference functional is CARA, the optimal unconstrained allocations may lie outside the feasible interval and therefore the optional constrained allocations may lie on the bounds. It is clear from the equations above, that when the optimal allocation x^* is either zero or m, its distribution is degenerated at the bounds. Implicity, this specification implies that the subject never make mistakes at the bounds. This is a quite strong assumption to make and in order to cover similar possibilities, we propose two additional specifications.

3.3.2 Specification 2 - CARA, Beta with Bias

In this specification, we assume that the random variables follow a Beta distribution with an additional *bias* parameter. More concretely, we define the following variables:

$$x_i' = \frac{bm}{3} + (1 - b_i)x_i^* \tag{3.13}$$

where b_i stands for the *bias* parameter and we replace the x_i^* from specification 1 with 3.13. When $b_i = 0$ there is no bias. When $b_i > 0$ there is bias and its magnitude depends on the value of the optimal allocation x_i^*. When $x_i^* = m/3$ then this bias goes to zero, which means that as one deviates from equal division, the bias increases. Applying this model of stochastic specification, eliminates the problems of degeneration that appear when one applies boundary portfolios. More specifically, if a portfolio contains an allocation on the boundaries, and provided that $b_i > 0$ then the distributions of the random variable are not degenerate. If this is the case, it is still possible to observe non-zero actual allocations when the optimal solution implies that the allocation should be zero.

3.3.3 Specification 3 - CARA, Two Betas

Next, we model the stochastic process of decision making in a slightly different way. In contrast to the two specifications presented above, where the stochastic specification remained the same for all values of the random variable (with bias or without), in this specification we assume that the error follows specification 1 when the allocation is within the bounds and when the allocation is on the bounds, either on zero or the total income (m), the error specification follows a different process. For this specification we introduce a new parameter d. It is assumed that when an allocation is equal to 0, the actual random variable x/m is distributed according to the beta distribution with parameters 1 and d. Consequently, the mean and the variance are given by the following formulas:

$$\text{Mean} = \frac{1}{1+d}$$

$$\text{Variance} = \frac{d}{(1+d)^2(d+2)}$$

When the allocation is equal to the total income m, then it is assumed that the variable is distributed with parameters d and 1 with the respective mean and variance to be given by the formulas:

$$\text{Mean} = \frac{d}{1+d}$$

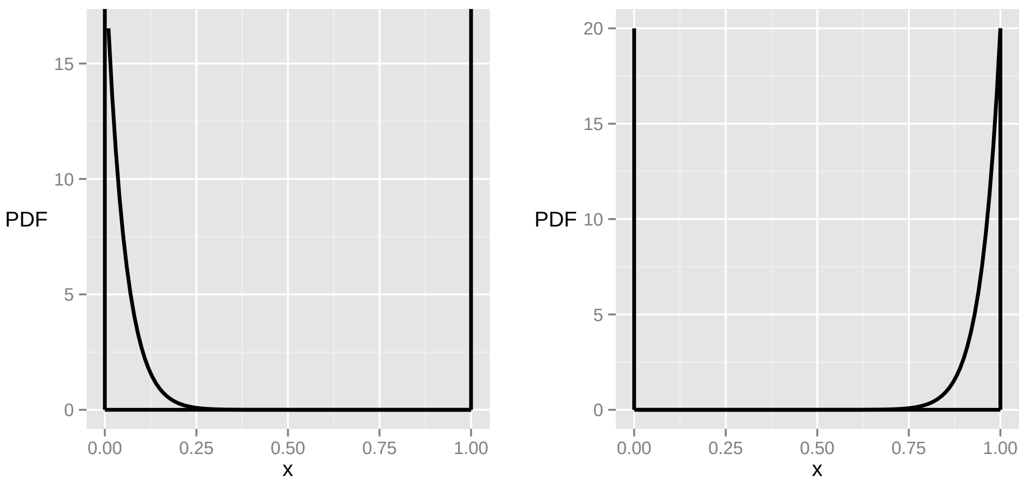

Figure 3.2 Beta Probability Distribution Function at the Bounds

$$\text{Variance} = \frac{d}{(1+d)^2(d+2)}$$

As it shown in Figure 3.2, when a random variable has a shape parameter α equal to 1 then its mode is at 0 and similarly, when the shape parameter β is 1 the mode is at 1, so the error gets larger as one deviates from the bounds. For the simulation, it is important to consider that the allocations should always sum up to the total income m. It is necessary to consider the various cases depending on the values of the optimal allocations x_i^*, x_j^*.

1. If all the allocations $\{x_i^*, x_j^*, x_k^*\}$ are positive, then the process is the same as in specification 1.

2. If one of the $\{x_i^*, x_j^*, x_k^*\}$ is equal to the total income m (the other two are equal to zero) one way of proceeding is to generate the actual allocation based on the beta shape parameters (1,d). This will necessarily provide a value lower than the total income. We then allocate equally the residual $(m - x_i)$ to the two remaining assets.

3. If one of the $\{x_i^*, x_j^*, x_k^*\}$ is equal to 0, we need to consider what holds for the two remaining allocations. If for example $x_i^* = 0$ then we assume that x_i is beta distributed with parameters (1,d). Then x_j is distributed according to beta with parameters $\alpha = \frac{x_j^*(s_j-1)}{m-x_i^*}, \beta = \frac{(m-x_i^*-x_j^*)(s_j-1)}{m-x_i^*}$).

Table 3.3, summarises the three specifications and their respective parameters.

Table 3.3 Specifications and Parameters

	Specification	Utility Function	Parameters
1	Beta	CRRA	r, s
2	Beta with Bias	CARA	r, s, b
3	Two Betas	CARA	r, s, d

3.4 A Simulation Study

In this section we report on extensive simulations inspired by the work of Wilcox (2008). More specifically, Wilcox suggests that sometimes, the stochastic specification may be more important than the preference functional (Wilcox (2008), p. 200):

> [...] when it comes to evaluating theories of discrete choice under risk, where many interesting inferences depend crucially on stochastic assumptions, average treatment effects alone are relatively uninformative.

The scope of Wilcox (2008), was to investigate and compare different error stories in the context of pairwise choice experiments on decision making under risk. The main objective was to test whether the error specification is a crucial decision variable for the experimenter. We follow a similar approach for our allocation problem and we report on simulation results, using data that that we have generated.

The task of the simulation is quite straightforward. Based on the three different specifications presented above, we cross-check if it is possible to identify the different stochastic specifications and to see what happens when we use the wrong specification when fitting preference functionals. We examine the estimation results from the 9 pairwise combinations of the three true error specifications and the three assumed-true specifications. The way to check this is by running the simulation, combining all the possible scenarios. This means, that after having expressed the closed form solutions for all the specifications, respecting the utility function constraints (CRRA for specification 1, CARA for 2 and 3), we use the appropriate estimation routine (for specifications 1, 2 and 3) and estimate all the different cases, based on simulated data that we have generated according to the respective specifications. This enables us to see if the inferences drawn are very different when the wrong specification is applied. As is expected, the results of a similar exercise, are quite sensitive to the underlying parameters. For this reason, we consider a number of different parameter sets, related to both the preferences of the subjects (risk aversion) and the underlying stochastic specification (precision s, bias b, d). For computational simplicity we adopt the following procedures. Initially, we normalise the value of the experimental income to 1. This means that the optimal allocations are directly expressed in proportions rather than in integers.[7] For a CRRA subject, as is shown in Table 3.1, a value of r greater or equal to 1 is indicating a person who is risk-loving or risk-neutral. If this is the case, the optimal decision in our task is to allocate everything on the asset that has the greater expected payoff and is defined as the asset for which it holds $p_s e_s >= \max\{p_i e_i, p_j e_j, p_k e_k\}$. As our task does not allow us to discriminate between risk-loving and risk-neutral persons, we obtain no information when the value of r is greater than one. Then, when the values of r is zero or negative the functional form of the utility function changes. The values for the risk aversion parameter that we consider here lie in the interval [0,1] as this interval covers a range of reasonably risk-averse subjects. Similarly, in the case of CARA subjects, the value of r can be positive, zero or negative leading to different functional forms. If r is zero, then the subject is risk-neutral and when the risk aversion parameter becomes negative, it indicates a risk-loving person. The range of values that we consider is [1,5]. It is important to notice the differences when r increases. An increase in the CRRA case indicates a reduction of risk aversion, while an increase in the CARA case indicates an increase of risk aversion. Although, there is no strict mapping between the values of r for CRRA and CARA, we have chosen the values in such a way that the highest value for r for a CRRA subject implies the same (low) amount of risk aversion (concavity) as the lowest value of r for the CARA subjects. In the same way, the lowest value of r for CRRA implies the same (high) amount of risk-aversion for CARA as the highest value of r in that case. As their is no full mapping between the risk aversion for the two cases, we suggest that the values for which the parameters of the two function captures the same levels of risk aversion, are those shown in Table 3.4. This also can be seen is

[7] As the Beta distribution requires the random variables to be strictly between 0 and 1, one can either divide the actual allocation by the income, which is always in the interval [0,1], or can apply the methodology we adopt.

Figures 3.3-3.6 which illustrate the utility functions for values $r = .1$ and $r = .9$ for the CRRA function and, $r = 1$ and $r = 5$ for the CARA.

Table 3.4 Similarity Between the Risk Aversion Coefficient

r **CRRA**	*r* **CARA**
0.9	1
0.7	2
0.5	3
0.3	4
0.1	5

Figure 3.3 CRRA, $r = .1$

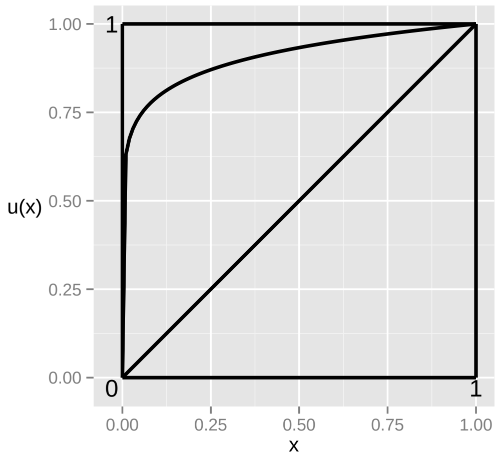

Figure 3.4 CRRA, $r = .9$

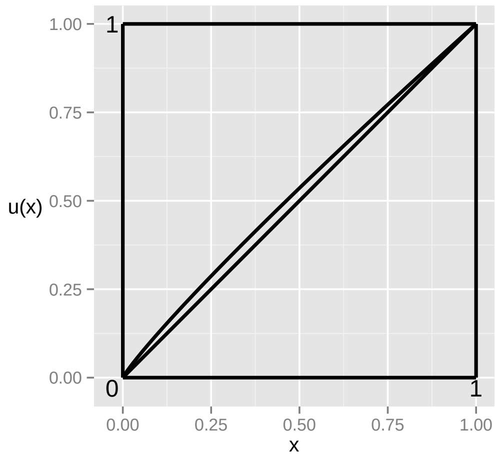

Figure 3.5 CARA, $r = 1$

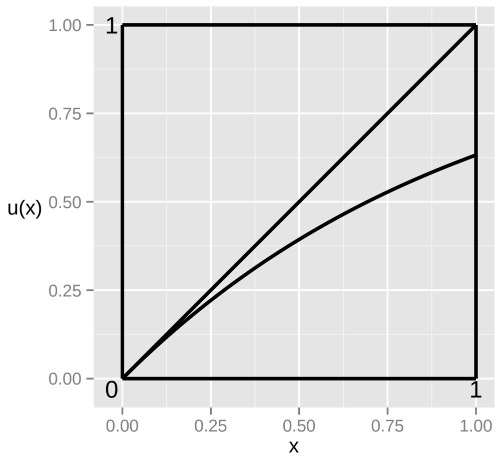

Figure 3.6 CARA, $r = 5$

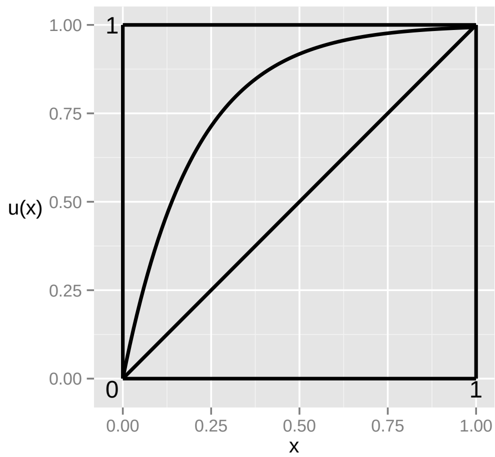

Table 3.5 lists the sets of parameters used for the simulation. There are 10 different sets that are divided in two different blocks. Within each block, the levels of risk aversion remain the same and what changes is the level of precision

(amount of noise in behaviour) with the first block being less noisy (more precise) than the second. Within each block, the values of r change in such a way to cover the whole feasible set of values they can take (from the lowest to the highest possible value).

Table 3.5 The Parameter Sets

Set	Risk Aversion Index r		Precision	Bias	Two Betas
	CRRA	CARA	s	b	d
1	0.9	1	50	0.05	40
2	0.7	2	50	0.05	40
3	0.5	3	50	0.05	40
4	0.3	4	50	0.05	40
5	0.1	5	50	0.05	40
6	0.9	1	25	0.1	20
7	0.7	2	25	0.1	20
8	0.5	3	25	0.1	20
9	0.3	4	25	0.1	20
10	0.1	5	25	0.1	20

We run the simulation on 72 different allocation problems[8] (combinations of probabilities and exchange rates). We assume that all states are treated equally, given their probabilities, and hence there is no psychological bias towards or away from particular states. The set of problems was chosen to span as much of the state space as possible, so we might expect a variety of behaviours from different subjects[9]. A total of 1000 simulations[10] was implemented. Tables 3.6-3.10, report the *means* and the *standard deviations* for all the parameters of interest (namely r, s, b, d and the value of the *maximised log-likelihood*.

Notice than no specification is nested inside any other. It may seem that specification 2 is nested inside specification 1 but they have different utility functions. The same holds for specifications 1 and 3. 2 and 3 are not nested either. The simulation exercise appears to be useful in two ways. Firstly, by running an extensive Monte Carlo simulation, it is possible to find out whether the problem under dis-

[8]We describe analytically, the way that these problems were chosen in the next section. The problems are also listed in Table 3.12.

[9]As is later indicated, one of the objectives of the simulation exercise, was also to enable us choosing a suitable problem set that will provide enough information on the decision process and we will be able to use it for an experimental test.

[10]The simulation program was written in the *R* programming language for statistical computing. The *R* Manuals, version 3.0.2. Available at: http://www.r-project.org/. The program is available upon request.

cussion is interesting and if it is proceeding to an empirical test by designing and running an economic experiment. The task seems to satisfy this role of simulation, as by assuming a different stochastic specification from the one that generated the data, results always to a lower log-likelihood predicted value. In addition, the simulation exercise is absolutely necessary to enhance the reliability of the analysis. Given that in a simulation the researcher is able to control and choose the values of the parameters of interest, by running extensive simulations and tests, it is ensured that the estimation program is working effectively and is also able to estimate and the underlying parameters.

We now present the results of this simulation. We begin by looking at Table 3.7, which reports the means and the standard deviations of the risk-aversion parameter r. Looking down the main diagonal of each block (where the true specification agrees with the estimated one) we see that everywhere the mean estimated parameter is close to the true value, as is expected. It is interesting that the CRRA estimated values seem to be closer to the true values than the CARA estimated ones[11].

However it is the off-diagonal elements that are most interesting and informative as these tell us about the dangers of *mis-specification*. One, should be very careful, since the risk aversion parameter for CRRA means something different from the risk aversion parameter for CARA. For example, look at Parameter set 1 (the first block) when the true specification is CRRA with r=0.9 and when specification 2 is used for estimation, the mean estimated value of r is 0.617.

In Table 3.8 we report the values of the precisions s. In this case, precisions are comparable. Again, along the main diagonal of each block the mean estimated precision is close to the true value of the precision, though the standard deviations are quite large. The latter is a consequence of the likelihood function being rather flat around its maximum, indicating that differences in precision do not make a big difference to behaviour. The off-diagonal elements, however, depart quite sharply from the true values. As a general rule, though it is not always the case, the estimated precision is less than the true precision. This is an interesting

[11] Nevertheless, it seems that the difference of the mean value of the risk aversion parameter r from its true value, is never significantly different at the 1% level.

result which suggests that mis-specification might lead to an under-estimation of the precision of the subjects. In Table 3.9, we report the results for the bias parameter (recall that this parameter applies only to Specification 2). When the true specification is Specification 1, the estimates of b are close to their true values but occasionally depart significantly from them. When the true value of b is zero, the estimated values are not significantly different from zero. This happens mostly in the parameter sets where the value of the risk aversion parameter indicates high levels of risk aversion. The estimates of the d parameter in Specification 3 appear in Table 3.10. (recall that this parameter applies only to Specification 3). It is interesting to note, however, that for the parameter sets with higher risk aversion, the estimated value significantly deviates compared to the sets with lower risk aversion. Finally, Table 3.6, reports the means and standard deviations of the minimised negative log-likelihoods. It seems that these log-likelihoods are comparable across specifications. What we had expected was that the entries down the main diagonal of each block would be the smallest in each row (remember that these numbers are the negative of the minimised log-likelihood) indicating that if one chooses between specifications on the basis of the maximised log-likelihoods, then one would always correctly identify the true specification. This is the case most of the times, apart from the parameter sets with high implied risk aversion (sets 4, 5, 9 and 10), where it seems that the values of the maximised log-likelihoods are quite close to each other, a fact that makes it difficult to distinguish between the specifications. This is something that was expected to happen, as in the case of high risk aversion, a decision maker, regardless of whether her utility is CRRA or CARA, will tend to equalise the payoffs for all states and consequently, will almost never have any incentive to make boundary allocations (the total income or nothing). Taking this for granted, it is reasonable to regard Specifications 2 and 3 as equivalent, as away from the bounds, the two specifications generate the same result. With the simulation results at hand, we run an incentivised experiment to test the three specifications in practice. In the remaining sections, we present the experimental design and report the results.

Table 3.6 Means and Standard Deviations of the *Maximised Log-Likelihoods*

Parameter Set 1

True	Estimated Specification 1	2	3
1	-137.35	-199.54	-278.37
	(7.046)	(8.942)	(19.155)
2	-401.44	-299.06	-409.04
	(9.297)	(10.7)	(17.454)
3	-451.16	-398.84	-307.12
	(12.233)	(18.107)	(11.405)

Parameter Set 2

True	Estimated Specification 1	2	3
1	-276.40	-348.63	-380.38
	(9.215)	(9.238)	(18.695)
2	-408.25	-348.72	-398.01
	(7.557)	(9.251)	(12.899)
3	-449.17	-413.89	-355.54
	(12.68)	(17.401)	(8.864)

Parameter Set 3

True	Estimated Specification 1	2	3
1	-350.78	-396.29	-427.31
	(8.986)	(7.176)	(9.623)
2	-471.01	-384.11	-389.52
	(6.782)	(8.485)	(7.977)
3	-470.25	-401.39	-391.31
	(9.207)	(11.898)	(9.938)

Parameter Set 4

True	Estimated Specification 1	2	3
1	-387.18	-410.13	-446.11
	(8.357)	(8.027)	(8.5)
2	-459.85	-412.37	-414.08
	(8.127)	(8.263)	(8.196)
3	-469.23	-406.56	-407.10
	(7.181)	(8.492)	(8.283)

Parameter Set 5

True	Estimated Specification 1	2	3
1	-405.66	-422.96	-445.28
	(8.232)	(8.845)	(6.511)
2	-467.01	-425.28	-425.89
	(8.789)	(8.062)	(8.394)
3	-457.62	-423.29	-423.14
	(8.795)	(8.215)	(8.246)

Parameter Set 6

True	Estimated Specification 1	2	3
1	-149.92	-195.05	-276.27
	(7.804)	(9.696)	(19.234)
2	-444.03	-375.04	-462.60
	(8.889)	(11.045)	(16.367)
3	-493.27	-455.34	-365.85
	(11.27)	(19.055)	(12.532)

Parameter Set 7

True	Estimated Specification 1	2	3
1	-307.61	-355.05	-387.37
	(10.585)	(10.35)	(19.203)
2	-439.55	-414.24	-448.13
	(9.154)	(9.36)	(10.358)
3	-488.69	-466.84	-412.97
	(11.858)	(14.301)	(9.91)

Parameter Set 8

True	Estimated Specification 1	2	3
1	-393.80	-418.98	-442.27
	(9.567)	(7.983)	(10.424)
2	-481.16	-443.08	-447.61
	(8.209)	(8.773)	(9.354)
3	-484.12	-455.66	-445.47
	(9.192)	(10.972)	(11.436)

Parameter Set 9

True	Estimated Specification 1	2	3
1	-433.44	-445.94	-470.02
	(8.513)	(8.558)	(8.426)
2	-488.11	-465.22	-466.56
	(8.894)	(8.218)	(8.295)
3	-492.15	-453.43	-454.00
	(7.729)	(8.182)	(8.588)

Parameter Set 10

True	Estimated Specification 1	2	3
1	-453.30	-463.64	-475.91
	(8.382)	(9.586)	(6.764)
2	-506.25	-476.62	-477.40
	(9.64)	(8.199)	(8.201)
3	-491.28	-471.38	-471.57
	(9.344)	(8.229)	(8.174)

Table 3.7 Means and Standard Deviations of the Estimated Value of the Risk-Aversion Parameter r

Parameter Set 1

True	Estimated Specification 1	2	3	True Value
1	0.900	0.617	1.846	0.900
	(0.003)	(0.114)	(1.5)	
2	0.606	1.025	1.275	1
	(0.015)	(0.015)	(0.485)	
3	0.471	1.012	1.002	1
	(0.039)	(0.234)	(0.036)	

Parameter Set 2

True	Estimated Specification 1	2	3	True Value
1	0.700	1.417	1.627	0.700
	(0.007)	(0.179)	(0.342)	
2	0.515	2.045	2.289	2
	(0.011)	(0.046)	(0.167)	
3	0.501	1.954	1.999	2
	(0.021)	(0.054)	(0.035)	

Parameter Set 3

True	Estimated Specification 1	2	3	True Value
1	0.500	2.153	2.685	0.500
	(0.011)	(0.123)	(0.253)	
2	0.345	3.075	3.241	3
	(0.016)	(0.059)	(0.075)	
3	0.365	2.885	3.019	3
	(0.019)	(0.072)	(0.075)	

Parameter Set 4

True	Estimated Specification 1	2	3	True Value
1	0.299	2.451	3.384	0.300
	(0.018)	(0.109)	(0.431)	
2	0.045	4.029	4.209	4
	(0.026)	(0.132)	(0.062)	
3	0.136	3.978	4.002	4
	(0.024)	(0.057)	(0.037)	

Parameter Set 5

True	Estimated Specification 1	2	3	True Value
1	0.100	2.822	4.467	0.100
	(0.027)	(0.327)	(0.073)	
2	0.010	4.979	5.272	5
	(0)	(0.306)	(0.147)	
3	0.010	4.876	5.013	5
	(0)	(0.229)	(0.125)	

Parameter Set 6

True	Estimated Specification 1	2	3	True Value
1	0.900	0.596	1.977	0.900
	(0.004)	(0.105)	(1.525)	
2	0.512	1.061	1.767	1
	(0.02)	(0.034)	(0.817)	
3	0.316	1.155	1.006	1
	(0.052)	(0.764)	(0.063)	

Parameter Set 7

True	Estimated Specification 1	2	3	True Value
1	0.700	1.413	1.643	0.700
	(0.011)	(0.196)	(0.383)	
2	0.431	2.109	2.646	2
	(0.019)	(0.077)	(0.27)	
3	0.414	1.903	1.999	2
	(0.031)	(0.129)	(0.063)	

Parameter Set 8

True	Estimated Specification 1	2	3	True Value
1	0.499	2.142	2.738	0.500
	(0.017)	(0.144)	(0.307)	
2	0.257	3.156	3.493	3
	(0.028)	(0.103)	(0.149)	
3	0.309	2.797	3.048	3
	(0.03)	(0.126)	(0.152)	

Parameter Set 9

True	Estimated Specification 1	2	3	True Value
1	0.299	2.460	3.525	0.300
	(0.027)	(0.149)	(0.485)	
2	0.012	4.063	4.445	4
	(0.008)	(0.23)	(0.11)	
3	0.138	3.954	4.000	4
	(0.035)	(0.104)	(0.056)	

Parameter Set 10

True	Estimated Specification 1	2	3	True Value
1	0.098	3.006	4.471	0.100
	(0.038)	(0.57)	(0.123)	
2	0.010	5.031	5.586	5
	(0)	(0.448)	(0.249)	
3	0.010	4.845	5.007	5
	(0)	(0.323)	(0.17)	

Table 3.8 Means and Standard Deviations of the Estimated Value of the Precision Parameter *s*

Parameter Set 1

	Estimated Specification 1	2	3	True Value
1	54.80 (10.224)	5.52 (1.287)	3.69 (2.126)	50
2	12.12 (1.324)	53.64 (7.468)	41.09 (16.714)	50
3	4.31 (0.474)	9.87 (2.985)	50.38 (10.825)	50

Parameter Set 2

True	Estimated Specification 1	2	3	True Value
1	54.32 (7.68)	15.78 (2.349)	15.78 (6.753)	50
2	21.53 (2.124)	52.67 (6.515)	41.35 (12.196)	50
3	11.68 (1.601)	20.15 (4.275)	51.89 (7.218)	50

Parameter Set 3

True	Estimated Specification 1	2	3	True Value
1	53.22 (6.612)	26.90 (2.851)	19.43 (4.276)	50
2	16.34 (1.434)	52.54 (6.524)	50.04 (6.014)	50
3	15.77 (1.944)	40.19 (6.877)	50.96 (7.59)	50

Parameter Set 4

True	Estimated Specification 1	2	3	True Value
1	52.15 (6.023)	37.73 (4.556)	25.52 (4.493)	50
2	27.39 (2.974)	52.25 (6.648)	50.99 (5.851)	50
3	22.54 (2.096)	51.54 (6.31)	51.72 (7.524)	50

Parameter Set 5

True	Estimated Specification 1	2	3	True Value
1	51.70 (5.852)	40.95 (5.529)	29.88 (3.085)	50
2	31.11 (3.236)	51.63 (5.827)	51.32 (5.991)	50
3	33.61 (3.592)	51.24 (5.967)	51.44 (5.962)	50

Parameter Set 6

True	Estimated Specification 1	2	3	True Value
1	27.12 (5.709)	4.95 (1.059)	3.27 (1.942)	25
2	10.20 (0.964)	25.71 (3.151)	17.30 (9.442)	25
3	3.51 (0.283)	5.96 (1.575)	25.65 (6.379)	25

Parameter Set 7

True	Estimated Specification 1	2	3	True Value
1	26.53 (3.638)	11.86 (1.852)	11.52 (4.105)	25
2	18.22 (2.025)	25.93 (3.011)	19.31 (4.767)	25
3	8.05 (0.923)	10.91 (1.757)	25.58 (3.896)	25

Parameter Set 8

True	Estimated Specification 1	2	3	True Value
1	26.13 (3.149)	17.70 (1.991)	13.66 (2.83)	25
2	15.72 (1.609)	25.98 (3.335)	24.86 (3.242)	25
3	13.31 (1.515)	19.45 (3.141)	25.00 (4.408)	25

Parameter Set 9

True	Estimated Specification 1	2	3	True Value
1	25.93 (2.856)	21.77 (2.582)	16.75 (2.756)	25
2	19.39 (2.127)	25.84 (3.079)	25.48 (2.887)	25
3	15.76 (1.55)	25.60 (2.886)	25.56 (3.672)	25

Parameter Set 10

True	Estimated Specification 1	2	3	True Value
1	25.76 (2.936)	22.41 (3.015)	18.92 (2.036)	25
2	18.65 (2.004)	25.63 (2.998)	25.46 (2.863)	25
3	20.62 (2.304)	25.63 (2.945)	25.61 (2.933)	25

Table 3.9 Means and Standard Deviations of the Estimated Value of the Bias Parameter *b*

True	Parameter Set 1 Estimated Specification 2	True Value	True	Parameter Set 6 Estimated Specification 2	True Value
1	0.01 (0.004)	0	1	0.01 (0.004)	0
2	0.04 (0.003)	0.05	2	0.08 (0.007)	0.1
3	0.07 (0.022)	-	3	0.13 (0.039)	-

True	Parameter Set 2 Estimated Specification 2	True Value	True	Parameter Set 7 Estimated Specification 2	True Value
1	0.03 (0.005)	0	1	0.03 (0.006)	0
2	0.04 (0.004)	0.05	2	0.07 (0.008)	0.1
3	0.05 (0.008)	-	3	0.09 (0.016)	-

True	Parameter Set 3 Estimated Specification 2	True Value	True	Parameter Set 8 Estimated Specification 2	True Value
1	0.07 (0.008)	0	1	0.07 (0.01)	0
2	0.03 (0.005)	0.05	2	0.07 (0.012)	0.1
3	0.04 (0.009)	-	3	0.07 (0.017)	-

True	Parameter Set 4 Estimated Specification 2	True Value	True	Parameter Set 9 Estimated Specification 2	True Value
1	0.13 (0.008)	0	1	0.13 (0.012)	0
2	0.04 (0.028)	0.05	2	0.08 (0.048)	0.1
3	0.00 (0.01)	-	3	0.01 (0.015)	-

True	Parameter Set 5 Estimated Specification 2	True Value	True	Parameter Set 10 Estimated Specification 2	True Value
1	0.18 (0.031)	0	1	0.16 (0.057)	0
2	0.05 (0.052)	0.05	2	0.10 (0.069)	0.1
3	0.03 (0.038)	-	3	0.03 (0.051)	-

Table 3.10 Means and Standard Deviations of the Estimated Value of the Second Beta Parameter *d*

True	Parameter Set 1 Estimated Specification 3	True Value	True	Parameter Set 6 Estimated Specification 3	True Value
1	47.70 (18.33)	-	1	46.30 (18.178)	-
2	11.66 (11.276)	-	2	13.27 (12.976)	-
3	41.46 (6.104)	40	3	21.00 (3.211)	20
True	**Parameter Set 2** **Estimated Specification** 3	**True Value**	**True**	**Parameter Set 7** **Estimated Specification** 3	**True Value**
1	51.68 (10.581)	-	1	52.72 (10.433)	-
2	30.22 (15.512)	-	2	25.59 (9.874)	-
3	41.75 (6.877)	40	3	20.82 (3.441)	20
True	**Parameter Set 3** **Estimated Specification** 3	**True Value**	**True**	**Parameter Set 8** **Estimated Specification** 3	**True Value**
1	39.33 (11.843)	-	1	41.63 (11.757)	-
2	55.22 (6.397)	-	2	31.44 (8.817)	-
3	42.10 (9.41)	40	3	21.53 (6.63)	20
True	**Parameter Set 4** **Estimated Specification** 3	**True Value**	**True**	**Parameter Set 9** **Estimated Specification** 3	**True Value**
1	21.29 (9.217)	-	1	22.71 (11.998)	-
2	30.90 (17.615)	-	2	28.96 (17.225)	-
3	34.36 (19.258)	40	3	36.90 (19.074)	20
True	**Parameter Set 5** **Estimated Specification** 3	**True Value**	**True**	**Parameter Set 10** **Estimated Specification** 3	**True Value**
1	29.96 (17.026)	-	1	30.16 (17.763)	-
2	29.29 (17.337)	-	2	30.96 (17.475)	-
3	30.00 (16.903)	40	3	30.06 (17.194)	20

3.5 The Experimental Procedure

The task that the subjects were presented with was the same as the one we discussed earlier. There were three different assets, represented by table tennis balls of three different colours (Red, Green and Blue). In total there were 10 balls. For each problem, the subjects were endowed with an amount of 100 tokens (which remained the same during the experiment) and were asked to allocate them to the three assets. In addition, they were told the exchange rates between tokens and British sterling, as well as the composition of the bag (the probabilities) if this problem was to be played for real. An incentive compatible mechanism was applied by picking at random one of the 72 problems and physically playing it out. So the subject filled the bag according to the probabilities of the chosen problem, shook the bag and then picked a ball at random, which defined the winning colour. The payment was defined by the product between the number of tokens that have been allocated to the winning colour and the respective exchange rate. The subjects were paid immediately after the end of the experiment in cash and privately. The experiment was conducted at the Centre for Experimental Economics (EXEC) at the University of York, with a total of 63 participants (35 females, 55%) in May-June 2014 and it lasted less than 45 minutes. There was a minimum period of 15 seconds where subjects could not submit their choices and they had maximum 1 minute per question to make their allocations. Each subject was allocated to an isolated terminal and the participants were not able to communicate with each other. The subjects were recruited using the *hroot* (Hamburg Registration and Organization Online Tool, Bock et al. (2012)) from a standard student population pool that included both undergraduate and postgraduate students. The fields of studies were diverse. Written instructions were provided (see appendix D) followed by a slide-presentation, where the decision task as well as the experimental framework were explained. Participants had the opportunity to ask clarifications and then they could start the experiment and proceed at their own pace, subject to the available minimum and maximum time to respond. The average payment was £14.62, with the maximum payment being £27.10 and the minimum £2.70 (st. dev. 4.21).

Table 3.11 Summary of the Sessions

Session	1	2	3	4	5	6	Overall
N	10	11	10	10	12	10	63
Av. Earn.	15.0	14.1	15.1	15.3	13.2	14.96	14.6

3.5.1 The Experimental Framework

The experiment was computerised and the experimental software was developed in *Python*[12]. For the allocations, an innovative graphical representation was implemented. The allocation space was represented by a simplex (an equilateral triangle) where each point inside the triangle corresponded to an allocation to the three assets, which was a function of the distances between the cursor and the sides of the triangle. Representing each asset at one of the vertices, choosing an allocation at a specific vertex meant allocating all the income only to one asset. Moving the mouse pointer, the subjects were able to choose their preferred allocation. In addition, a bar-chart on the screen provided information to the subjects. The width of the bars was proportional to the implied probability for each colour, while the height was showing the implied payment if this state of the world occured. The exchange rates were expressed in pennies. A timer informed the subjects of their remaining time to respond. In addition, for each problem, there was a graphical representation of the composition of the bag as a visual aid. Figures 3.7 and 3.8 show screenshots from the experimental interface.

3.5.2 The Problems

Table 3.12 presents the full set of 72 problems on which participants were asked to make their allocations. As was described before, the problems were carefully chosen in such a way that it would be possible to distinguish between the different specifications. The choice of the probabilities for the three colours was decided to include all the possible combination of probabilities that ensures that there is always at least 1 ball of each colour. Practically this means, that the minimum probability for a state is equal to .1 and the maximum equal to .8. With this constraint as given, we consider all the possible combinations that ensure that the sum of all the balls is equal to 10 (or in probability terms, the chances

[12]Python Software Foundation. Python Language Reference, version 2.7. Available at http://www.python.org. The software is available upon request.

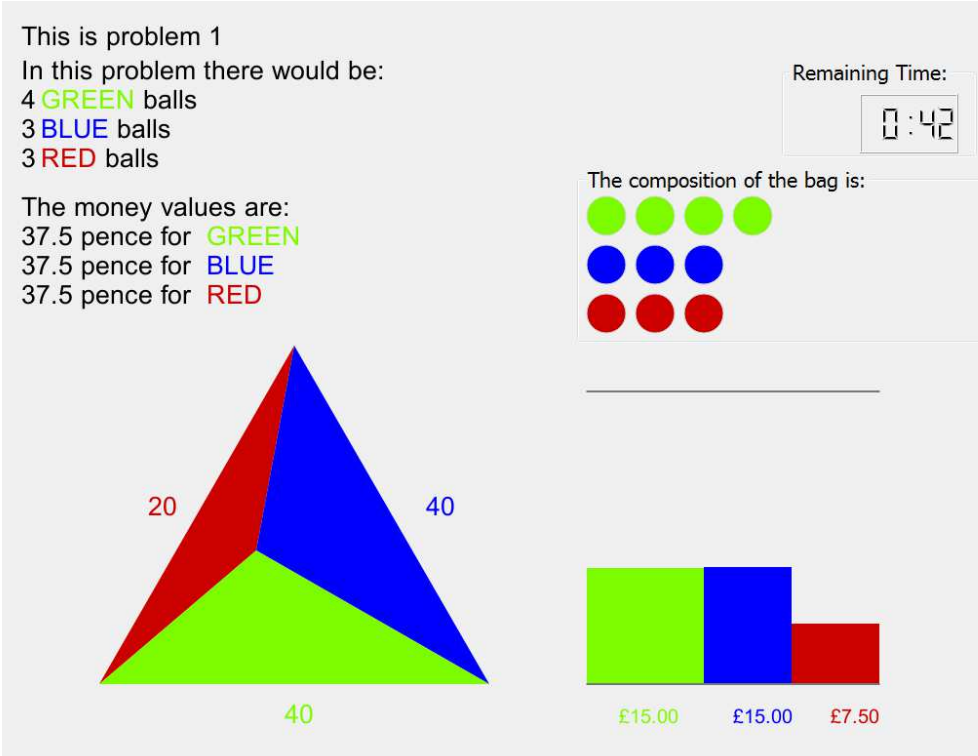

Figure 3.7 Screenshot 1 of the Experimental Framework

add up to 1). Then, we assume that no matter what the colour of the ball is, the symmetric problems are treated in a similar way. For example that the set of probabilities (.2,.3,.5) is observationally equivalent to the set (.5,.3,.2) and it will not provide any additional valuable information to consider both, assuming always that the assets are treated in a symmetric way regarding their characteristics. For instance, with the exchange rate fixed to 1, the only important criterion is the probability of occurring and not any of the irrelevant characteristics (e.g. the colour). Eliminating all the combinations that have the above characteristics, it leaves 8 possible combinations. Then, we assume three different levels of exchange rates. In the simulation, we considered a low value, an intermediate and a high value and more specifically the values .75, 1 and 1.25. These exchange rates are in terms of money values, so if one allocated 10 experimental income units to an asset that has exchange rate equal to 1.25, then if the respective state is also the one that actually occurs, these 10 units are transformed to £12.50. Again, considering all the possible combinations of the exchange rate and eliminating duplicates, leaves 9 combinations. Combining probabilities and exchange rates, we obtain the total of the 72 problems.

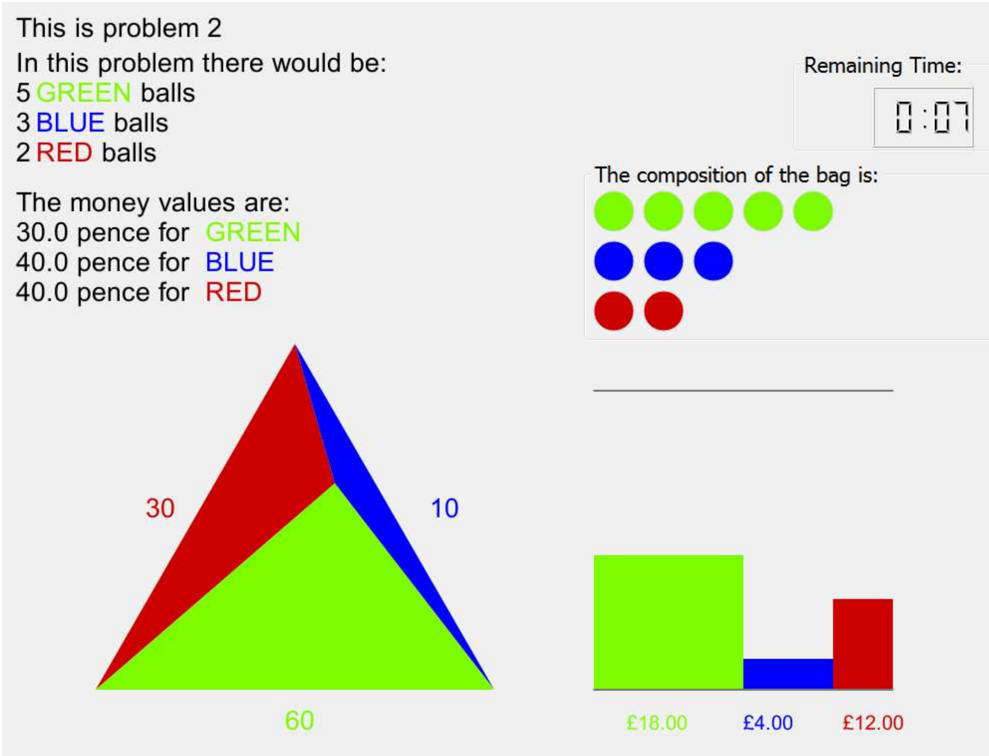

Figure 3.8 Screenshot 2 of the Experimental Framework

The problems were chosen to satisfy two criteria. First, the set of questions were designed, after extensive Monte Carlo simulations, in such a way that the produced data-set should be rich enough to identify the different specifications (the simulation is discussed in section 3.4). The second criterion was to choose the different problems in such a way that it was ensured for a risk-neutral decision maker[13] will always expect to win on average £15. The exchange rates were re-scaled, based on the formula $s = \frac{15}{\max\{e_i p_i, e_j p_j, e_k p_k\}}$ which guarantees the latter. The order of the problems was randomised for each subject in order to eliminate possible *order effects*.

[13]It is reminded that a risk neutral decision maker, allocates the total amount of the income to the asset with the highest expected payoff. Since the income is normalised to 1, this happens on the asset that has the highest product between the probability of occurring and the exchange rate.

Table 3.12 Problems of the Experiment

Problem	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
p_i	0.4	0.4	0.5	0.5	0.6	0.6	0.7	0.8	0.4	0.4	0.5	0.5	0.6	0.6	0.7	0.8	0.4	0.4
p_j	0.3	0.4	0.3	0.4	0.2	0.3	0.2	0.1	0.3	0.4	0.3	0.4	0.2	0.3	0.2	0.1	0.3	0.4
p_k	0.3	0.2	0.2	0.1	0.2	0.1	0.1	0.1	0.3	0.2	0.2	0.1	0.2	0.1	0.1	0.1	0.3	0.2
e_i	37.5	37.5	30	30	25	25	21.4	18.8	37.5	37.5	30	30	25	25	21.4	18.8	30	37.5
e_j	37.5	37.5	30	30	25	25	21.4	18.8	37.5	37.5	30	30	25	25	21.4	18.8	30	37.5
e_k	37.5	37.5	30	30	25	25	21.4	18.8	50	50	40	40	33.3	33.3	28.6	25	50	62.5

Problem	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36
p_i	0.5	0.5	0.6	0.6	0.7	0.8	0.4	0.4	0.5	0.5	0.6	0.6	0.7	0.8	0.4	0.4	0.5	0.5
p_j	0.3	0.4	0.2	0.3	0.2	0.1	0.3	0.4	0.3	0.4	0.2	0.3	0.2	0.1	0.3	0.4	0.3	0.4
p_k	0.2	0.1	0.2	0.1	0.1	0.1	0.3	0.2	0.2	0.1	0.2	0.1	0.1	0.1	0.3	0.2	0.2	0.1
e_i	30	30	25	25	21.4	18.8	37.5	28.1	30	28.1	25	25	21.4	18.8	37.5	28.1	30	28.1
e_j	30	30	25	25	21.4	18.8	50	37.5	40	37.5	33.3	33.3	28.6	25	50	37.5	40	37.5
e_k	50	50	41.7	41.7	35.7	31.3	37.5	28.1	30	28.1	25	25	21.4	18.8	50	37.5	40	37.5

Problem	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54
p_i	0.6	0.6	0.7	0.8	0.4	0.4	0.5	0.5	0.6	0.6	0.7	0.8	0.4	0.4	0.5	0.5	0.6	0.6
p_j	0.2	0.3	0.2	0.1	0.3	0.4	0.3	0.4	0.2	0.3	0.2	0.1	0.3	0.4	0.3	0.4	0.2	0.3
p_k	0.2	0.1	0.1	0.1	0.3	0.2	0.2	0.1	0.2	0.1	0.1	0.1	0.3	0.2	0.2	0.1	0.2	0.1
e_i	25	25	21.4	18.8	30	28.1	30	28.1	25	25	21.4	18.8	30	22.5	30	22.5	25	25
e_j	33.3	33.3	28.6	25	40	37.5	40	37.5	33.3	33.3	28.6	25	50	37.5	50	37.5	41.7	41.7
e_k	33.3	33.3	28.6	25	50	46.9	50	46.9	41.7	41.7	35.7	31.3	30	22.5	30	22.5	25	25

Problem	55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72
p_i	0.7	0.8	0.4	0.4	0.5	0.5	0.6	0.6	0.7	0.8	0.4	0.4	0.5	0.5	0.6	0.6	0.7	0.8
p_j	0.2	0.1	0.3	0.4	0.3	0.4	0.2	0.3	0.2	0.1	0.3	0.4	0.3	0.4	0.2	0.3	0.2	0.1
p_k	0.1	0.1	0.3	0.2	0.2	0.1	0.2	0.1	0.1	0.1	0.3	0.2	0.2	0.1	0.2	0.1	0.1	0.1
e_i	21.4	18.8	30	22.5	30	22.5	25	25	21.4	18.8	30	22.5	30	22.5	25	25	21.4	18.8
e_j	35.7	31.3	50	37.5	50	37.5	41.7	41.7	35.7	31.3	50	37.5	50	37.5	41.7	41.7	35.7	31.3
e_k	21.4	18.8	40	30	40	30	33.3	33.3	28.6	25	50	37.5	50	37.5	41.7	41.7	35.7	31.3

3.6 Econometric Specification

For each specification, the marginal contribution to the log-likelihood, depends on the utility function assumed, as well as the form of the stochastic term. In specification 1, it is quite straightforward to form the log-likelihood function. Since the CRRA utility function is assumed, it is not possible to observe zero or negative optimal allocations. The contribution is formed assuming that the actual allocation is centered to the optimal allocation. In specification 2, it is assumed that subjects are represented by a CARA utility function. If this is a case, it is possible to see zero or even negative allocations if one solves the unconstrained problem. Based on the algorithm for the optimal constrained allocations, the marginal contribution to the log-likelihood is formed, centered to this optimal constrained allocation.

In specification 3, as it again contains allocations to the bounds, it is necessary to take into account the several different cases. The simplest case is the one where the portfolio consists of non-zero allocations. This simplifies things as the form of the distribution in this case is identical to the one of specification 1. The contributions to the log-likelihood are based on the actual and optimal allocations to assets i and j and their respective shape parameters. In order to form the log-likelihood function, we follow the same procedure as in chapter 2. Then, we need to consider the extreme cases, that is to say the cases where all the income is allocated to one asset, and then the cases where zero is allocated to the asset. This part is quite important as it defines the way we treat the marginal contributions (the way we assume noise is added to the data) Focusing first on the event where all income is allocated to one of the three outcomes. Let for example the case $x_i^* = m$ (which means that $x_j^* = x_k^* = 0$. Following the formulation presented in section 3.3.3, the random variable x_i is distributed according to the Beta distribution with shape parameters $\alpha = d, \beta = 1$. Then, we assume that in a similar way x_j is Beta distributed with symmetric shape parameters $\alpha = 1, \beta = d$. Finally, the third variable is assumed to be the residual. In this stage, one needs to be careful as which variables contribute to the log-likelihood function. Table 3.13 summarises all the possible cases. The estimation of the parameters was done by applying *Constrained Maximum Likelihood techniques*[14].

[14]The estimation program was written using in the *R* programming language for statistical computing. The program is available upon request.

Table 3.13 Shape Parameters - Specification 3

Case	x_i^*	x_j^*	x_k^*	α_{x_i}	β_{x_i}	α_{x_j}	β_{x_j}	α_{x_k}	β_{x_k}
1	$x_i^* > 0$	$x_j^* > 0$	$x_k^* > 0$	$\frac{x_i^*}{m}(s_i - 1)$	$(1 - \frac{x_i^*}{m})(s_i - 1)$	$\frac{x_j^*}{m-x_j^*}(s_j - 1)$	$(1 - \frac{x_j^*}{m-x_j^*})(s_j - 1)$	-	-
2	$x_i^* = m$	$x_j^* = 0$	$x_k^* = 0$	d	1	1	d	-	-
3	$x_i^* = 0$	$x_j^* = m$	$x_k^* = 0$	-	-	d	1	1	d
4	$x_i^* = 0$	$x_j^* = 0$	$x_k^* = m$	1	d	-	-	d	1
5	$x_i^* = 0$	$x_j^* > 0$	$x_k^* > 0$	-	-	$\frac{x_j^*}{m}(s_j - 1)$	$(1 - \frac{x_j^*}{m})(s_j - 1)$	$\frac{x_k^*}{m-x_k^*}(s_k - 1)$	$(1 - \frac{x_k^*}{m-x_k^*})(s_k - 1)$
6	$x_i^* > 0$	$x_j^* = 0$	$x_k^* > 0$	$\frac{x_i^*}{m}(s_i - 1)$	$(1 - \frac{x_i^*}{m})(s_i - 1)$	-	-	$\frac{x_k^*}{m-x_k^*}(s_k - 1)$	$(1 - \frac{x_k^*}{m-x_k^*})(s_k - 1)$
7	$x_i^* > 0$	$x_j^* > 0$	$x_k^* = 0$	$\frac{x_i^*}{m}(s_i - 1)$	$(1 - \frac{x_i^*}{m})(s_i - 1)$	$\frac{x_j^*}{m-x_j^*}(s_j - 1)$	$(1 - \frac{x_j^*}{m-x_j^*})(s_j - 1)$	-	-

3.7 Results and Discussion

In Table 3.14, we present the mean values of the estimated parameters for the three specifications (in parenthesis the standard deviations).

Table 3.14 Estimates of the Parameters

Specification	r	s	b	d	LL
1	-2.30	38.79	-	-	-459.15
	(3.99)	(29.61)	-	-	(60.77)
2	19.08	31.05	0.12	-	-465.96
	(6.22)	(17.95)	(0.07)	-	(54.16)
3	20.88	28.64	-	30.40	-474.17
	(5.04)	(17.63)	-	(18.55)	(54.24)

Notice that the values for the risk-aversion parameter. When running the simulation we were expecting a positive value for the risk aversion parameter in the case of CRRA and a value in the range [1,7] for the CARA. Looking superficially at the data, it is possible for one to observe that there are not many boundary allocations (there are few zero allocations) and this happens for the majority of the subjects and the majority of the problems. What one observes in Table 3.14 is that the estimated values of the risk aversion parameters, are remarkably high, indicating high levels of risk aversion. This raises the question of whether allocation problems make the subjects to behave in a more risk averse way, compared to the pairwise choice tasks or the multiple price lists. The latter requires further empirical investigation.

Parameters s and d are the precision parameter of the stochastic specifications. The larger their value, the less noise exists in the data. On average, specification 1 obtains the highest value of this precision (38.79) while for specifications 2 and 3 is high as well and significantly different than zero. Parameter b captures the bias when there are boundary allocations. Its value is inside the expected range as should have been zero if there were not any allocations 0 or m, which was not the case in our experiment as boundary allocations were observed for some of the subjects and problems.

Let us now focus on the goodness of fit. If one focuses only on the mean value of the maximised log-likelihood, she can infer that Specification 1 fits the data better than the other two. But in order to be able to make a direct comparison and since the three specifications differ in the number of the parameters that they require

(2 in Specification 1 and 3 in Specifications 2 and 3), one needs to correct for the degrees of freedom. Table 3.15, presents the mean values (standard deviations in parenthesis) for the corrected likelihood criteria namely, the Bayesian (BIC), the Akaike (AIC) and the corrected-Akaike (AIC_c) Information Criterion[15]. The criteria have been calculated for the 72 observations and the respective number of parameters for each specification. A lower value for the corrected likelihoods, implies a better fit.

Table 3.15 Corrected Log-Likelihoods

Specification	LL	BIC	AIC	AIC_C
1	-459.153	939.98	922.31	922.48
	(60.77)	(120.57)	(120.57)	(120.57)
2	-465.96	966.28	937.93	938.28
	(54.16)	(107.46)	(107.46)	(107.46)
3	-474.17	982.7	954.34	954.7
	(54.24)	(107.61)	(107.61)	(107.61)

It is obvious that no matter which criterion we are using, on average, the means in Specification 1 seem to be consistently higher (lower for the corrected likelihoods) than the other two, fact that indicated a better fitting. Let us now see what happens at the individual level. Using the corrected (for degree of freedom log-likelihood, BIC) so that we can compare within specifications, we rank the Specifications from best to worst (lower to higher value of the BIC). Table 3.16, shows the cumulative percentage in each ranking position. Regarding the specification with the best fitting, Specification 1 is ranked first for 82% of the subjects, is ranked first or second for 90 % of the subjects and so on. It is obvious that Specifications 1 and 2, are ranked either first or second for the majority of the participants.

Table 3.16 Rankings Based on the BIC

Specification	1	1-2	1-3
1	82	90	100
2	16	70	100
3	2	40	100

All values represent cumulative percentages

From the results above, it is inferred that Specifications 1 and 2 fit better that

[15]These criteria are analytically presented in chapter 2.

the third, and also that Specification 1 is relatively better that 2.

We now turn to the *statistical significance* of these results. As mentioned before, the three Specifications are not *nested* inside each other. In this case, a *Clarke* test is appropriate to test for significance (Clarke (2007)). The Clarke test is a distribution-free test used for comparing non-nested models. The null hypothesis behind this test is:

$$H_0 : P(L_{1i} - L_{2i} > 0) = .5$$

where L_{1i}, L_{2i} are the individual log-likelihoods for each of the models. The individual log-likelihood, is calculated for each of the 72 problems, based on the values of the estimated parameters. The test statistic then is:

$$T = \sum_{i}^{72} I(L_{1i} - L_{2i})$$

where I is an indicator function that takes the following values:

$$I = \begin{cases} 1 & L_{1i} - L_{2i} > 0 \\ 0 & L_{1i} - L_{2i} \leq 0 \end{cases}$$

Then the test statistic is based on a Binomial test assuming that $T \sim \text{Bin}(72, .5)$. The null-hypothesis is rejected when T=44 at 5% significance level and T=46 at 1%. Table 3.17, reports the results of a Clarke test. We cross-check all specifications. For example, in the first row, we test whether Specification 1 is significantly better than 2 and 3. This statistical test is telling us that Specifications 1 and 2 perform better

Table 3.17 Significance Test for Superiority of Specifications

	1	2	3
1	n.a.	39	52
2	17	n.a.	42
3	14	11	n.a

Clarke Test at 5% Significance Level. All values represent percentages.

than the third. Indeed, Specification 1 is significantly better than 2 for 39% of the subjects and is better for 52% of the participants compared to 3. For Specification 2, it seems that it is significantly better than 1 for 17% of the subjects and better than 3 for 42%.

3.8 Conclusion

As allocation questions are gradually becoming popular and are used more frequently in experimental studies, it is of paramount importance to develop appropriate ways of modelling these error stories. In this chapter we presented three specifications of modelling stochastic choice in 3-way allocation problems. We included combinations of both CRRA and CARA utility functions with the Beta-distribution which was assumed to represent the stochastic process that underlies the decisions of the subjects. We ran an extended simulation in order to obtain some intuition of what happens when the researcher assumes the wrong specification when the data are generated by a different one. This has effects to both the magnitude of the values of the estimated parameters, but also to the model selection process, as one needs to test for all the possible specifications. In addition, we present an economic experiment that we specifically design in order to obtain data that could allow us to test for the difference in the stochastic specification. We found that the CRRA specification fits best, even though it has fewer degree of freedom. That is good news, as using specification 1 means there is one parameter less that is needed to be estimated. Nevertheless, if one wants to use a CARA utility function, specification 2 that assumes the existence of a bias parameter can be used. The experiment showed that subjects exhibit high degrees of risk aversion. The latter means that boundary allocations are not met very frequently and consequently both CRRA and CARA could be used without the need to take into consideration boundary allocation. Further investigation is needed concerning the correlation between risk aversion and allocation problems or stating it differently, whether allocation question make subjects to behave in a more risk averse way compared to pairwise choice tasks.

Conclusion

In this section we conclude. We present the conclusions in the form of questions where the answers aim to present what is the main problem and the main research questions, what is the methodology applied, what are the main results of this study, what is the contribution of this study, what are the limitations and how this work can be extended.

What is the problem and why it is important?
How people update the prior beliefs that have been formed in an ambiguous environment is a question in the literature that has been partially or inadequately answered. The discussion orbits around dynamic decision making, or more precisely, sequential decision making. This is to be distinguished from the dynamic choice that involves long time period where the future is discounted. The sequential choice concerns the formation of prior beliefs, the reception of information and the updating of these priors. Discounting or impatience plays no role in this framework. In this thesis, the main objective is to provide some insights of how people update their initial beliefs that have been formed in an ambiguous environment, whether they are dynamically consistent or the history of events plays no role to the decisions and if they are dynamically inconsistent, how do they actually resolve this inconsistency. The latter raises several behavioural issues, since we assume that not all decision makers behave in the same way, and consequently we allow for various levels of sophistication in the decisions.

These questions are of significance if one takes into consideration that a considerable number of choices in economic life require this type of sequentiality. Savings decisions, financial decisions, investments, all are examples where both ambiguity and reception of information are present. The standard model in economic theory assumes that decisions are made based on the Expected Utility model and dynamic consistency is always satisfied. Experimental evidence shows that this is not the case. The question now is whether one wants to proceed in a *descriptive* or a *normative* way. If the objective is to predict behaviour, then assuming Expected Utility will lead to failures in capturing ambiguity aversion. Take for example public policy. If a new policy is to be applied and Expected Utility is assumed, if people's preference are non Expected Utility, the actions that the model predicts will be significantly different from the action that economic agents will actually take.

What do we already know?

Topics of decision making under ambiguity have been studied in experimental settings relatively early. Nevertheless, the majority of this work has been published during the last five years. This is mostly due to the lack of a theoretical framework, rich enough to capture ambiguity aversion on the one hand, and inadequacy of experimental methods and subsequently methods of analysis of the experimental data on the other. Empirical evidence converges to two main results, that the participants in economic experiments express attitudes towards ambiguity and that subjects are characterised by high degree of heterogeneity regarding these attitudes. This conclusion is drawn from experiments conducted in a static framework. When the problem is extended to its dynamic version, the evidence is not rich enough to allow conclusions to be safely inferred. The two available studies provide evidence of the violation of dynamic consistency. The drawback of these studies is that they are prone to criticism regarding the methodology adopted, as they include very few questions, one of them does not include any kind of incentives and none of them takes into consideration noise in the data.

How this work differs from what has been already done?

In order to provide answers to the questions that we aim to investigate, we designed economic experiments that incorporated features that would enable us to obtain data that would be appropriate to do so. As is highlighted by the results of Antoniou et al. (2013), failing to correct for the non-linearity of the utility function provides stronger behavioural support for the Bayesian rule. They conclude that in order to be able to capture ambiguity aversion, there are three extensions than need to be done, in theoretical, experimental and econometrics terms. Our experiments were designed in such a way that they would be able to capture all of the above. Our protocols, differ from what has been done in the literature in four distinct ways. Regarding the theoretical extensions, we include several different preference functionals that allow for ambiguity aversion. More specifically, we assume that the subjects in our experiment have preferences that can be explained either by the Subjective Expected Utility (SEU), the MaxMin Expected Utility (MEU), the Choquet Expected Utility (CEU) or a parsimonious version of the CEU, called the Source Choquet Expected Utility (SCEU). We aim to elicit beliefs by using an alternative method compared to the standard that has been used in the literature (Holt-Laury price lists and the Becker-DeGroot-Marschak mechanism). Using this kind of elicitation method, one is able to construct structural models that are capable of capturing attitudes towards ambiguity Then, the extensions in the experimental framework and the subsequent econometric analysis require tasks that will make the identification of different theoretical structures possible. Starting with the decision task, we use allocation problems contrary to the conventional method of pairwise choices. This kind of questions seems to provide informationally richer data-sets that render our objective possible. We create ambiguity in the lab using a Bingo Blower, a device that provides a transparent and non-manipulable representation of ambiguous events. In the analysis of the experimental data, we seriously take into consideration the effects that noise in the data has to the analysis and interpretation of the results. Due to the nature of the decision task, it is not possible to use standard techniques that are broadly applied (e.g. normality in errors). Consequently, different specifications are needed in order to proceed to the econometric estimation of the models. In one of the experiments we use a particular specification due to the flexibility that it allows. Nevertheless, as we recognize its limits, in the second experiment we propose three different specifications and we test them against each other in order to infer which is the best to be applied in similar frameworks. In addition, we ask subjects a considerably large set of questions (60 and 72) in order to reduce the levels of noise in the data and to obtain more robust results. Finally, the stochastic specification along with the parametrisation of the preference functionals and combined with the fact that the analysis is done on a subject level analysis, allows for the identification and classification of several different structures that aim to model decision making.

Are people the same?

We take as a point of departure the idea that people are different. This is not an unreasonable assumption to make. People differ in their tastes, their beliefs and the strategies applied to solve complex problems. Many of the experimental studies analyse the data assuming a *representative* economic agent and try to make inferences for the whole experimental population. This approach is adopted mostly for reasons of simplification. In both our experiments, we proceed by analysing the data in a subject level analysis. In the experiment in chapter 2 we define three different types of decision makers (resolute, naive and sophisticated) and we find that indeed subjects can be classified to different types. In the experiment in chapter 3, we assume three different specifications of modelling stochastic choice and we try to find which one fits the data best for each individual data-set. We find considerable heterogeneity which is an indication that analysing the data assuming only one type of decision makers may lead to totally different results.

Do people make mistakes?

We agree with the point of view that the stochastic part of decision making is quite powerful to explain behaviour and one should be extremely careful to take this into consideration when analysing experimental data. In chapter 2, we use an error specification that allows us to use the fewest possible parameters in the specification. This was done with the objective to reduce the levels of complexity in mind. In chapter 3 being motivated by the technical difficulties and constraints that the previous specification exhibits, we propose and test three different specifications. The results shows that when subjects are characterised by high degrees of risk aversion, the specification applied to chapter 2 is the one that best explains behaviour and as a consequence, using the simplest specification with the fewest parameters, will not have any significant effect on the analysis.

Are people dynamically consistent?

As is expected, the answer to this question is not quite definite. From our experiment, we find that although having non-EU preferences, a large proportion of the experimental population (almost 1/3) is dynamically consistent. This result is contrary to the mis-conception that a non-EU decision maker should obligatory be dynamically inconsistent. The results do not provide support to the standard SEU model. Instead, non-EU models (more specifically SCEU) seems to be able to fit well to the decisions of more than 30% of the subjects. People have non-EU preferences, but decide that in the second stage they want to keep their choices constant. On the other hand, there is an equivalent proportion of subjects, that behave in a quite dynamically inconsistent way. Not only do they have non-EU preferences, but they also behave naively, failing to take into consideration the further stages of the task and failing to plan ahead. The rest of the subjects' behaviour is explained by a mixture of non-EU preferences and dynamic consistency or consequentialism. All of the above are extremely important, when the behavioural implications of these inconsistencies need to be taken into consideration (e.g. reaction of agents to a change in the taxation regime) as there may be required some form of government intervention (nudge) in the presence of dynamic inconsistencies, or there may not in cases where people are consistent with their choices. The latter raises some issues on inconsistency and welfare which we discuss later.

What are the limitations of this approach?

It goes without saying that our results are heavily based on the assumptions we make regarding the decision process, the stochastic specification and the representation of tastes and beliefs. In addition, the results are a function of the way that ambiguity was represented in the lab. There is no direct way that one can generalise the results above. One of the limitations of the methodology we applied can be found in the number of the parameters. The preference functionals we tested varied from having 4 to 8 parameters, from the simplest specification of SEU to the more complicated and non-smooth (CEU). Estimating a large number of parameters using maximum likelihood techniques is a computationally intense task and the convergence of the optimisation problem is sensitive to many factors (e.g. starting parameters). Alternative ways could be applied in order to analyse the results (e.g. Houser et al. (2004)). Nevertheless, it is not clear whether these methods could eliminate the current difficulties or would generate additional hurdles. During the experiment in Chapter 2, there was the implicit assumption that the type of the decision maker remains unchanged throughout the duration of the experiment, as did the preferences towards ambiguity. This is a strong assumption that needs further testing. Another limitation is to be found to the limits of the decision task. As in both experiments we were requiring subjects to allocate their total income, we were not able to observe risk loving behaviour (we could observe boundary allocations, but on the bounds there is no way to observationally distinguish between the risk neutral and the risk loving decision maker). Modifications of the decision task, such as allowing for negative allocations, would solve this issue. But as was argued before, is not straightforward how the questions of the experiment can be efficiently designed (how to provide a large enough amount of money to the subjects). The representation of ambiguity plays a significant role for the interpretation of the results. The Bingo Blower provides an adequate source of ambiguity but there were a number of technical constraints that reduced the efficiency of the device. Recall, that during the experiment we were asking subjects to suppose that a scenario holds, and then make their decisions accordingly. We did so as we wanted to avoid any kind of learning effects that would distort the objective of the experiment. An alternative would be to draw balls for real during the experiment or even to play every problem for real and then rearranging the composition of the Bingo Blower. Something like this would require a considerable amount of time and a lot of manipulation to hide the Bingo Blower while refilling that would raise suspicion on the one hand and would make the subjects to feel bored or even become confused on the other. Inevitably we needed to proceed with the assumed scenarios that would be played out for real at the end of the experiment. Another solution to this, is to use computerised sources of ambiguity. While this solution is the most efficient in terms of presentation and application, it heavily suffers from suspicion issues, as subjects may always suspect that the software has been programmed to the benefit of the experimenters. The use of the Bingo Blower creates an additional problem regarding the number of periods that the task can be extended to. The optimal number of colours is 3. Having less than 3 colours renders the task trivial. To make the problem a 4-period problem, 4 different states of the world, thus 4 colours. The consequence of this is that due to the limited space of the device and the requirement of having lot of balls to create ambiguity, it will be extremely difficult for subjects to distinguish the different proportions of the balls and would probably confuse them.

What is the plan for future work?
As is expected, this line of research provides significant evidence to the literature of decision making under ambiguity while at the same time it opens the path for some important direct extensions. The experiments we run generated a number of research questions that are connected either directly to the issue of updating, or to further methodological issues. We can identify three possible routes that this research can be extended individual decision making, interacting decision making and methodology. On individual decision making, the direct extension is to add additional periods to the decision task. At each period, some of the states of the world will be excluded and the priors will be updated accordingly. This requires the extension of the decision trees and also the way that ambiguity is represented (we extend this later). Extending the experiment to multiple-periods, allows us to test two different issues. On the one hand, it is possible to design a life cycle savings problem in the spirit of Carbone and Infante (2014), where subjects will be asked to make decisions for longer horizons. This extension requires the elongation of the time-horizon and the relaxation of the constraint that the total income must be allocated in one-shot. Subjects may be able to save or borrow[16] money and thus, make decisions in an environment that better resembles the real world economic environment. On the other hand, extending the number of periods, will provide subjects the opportunity to learn something by a process of partially receiving information. The theoretical framework of learning under ambiguity has been extensively developed (Epstein and Schneider (2007), Epstein et al. (2010)) and few experiments test some of the available models (see literature review in chapter 1). Combining the two would provide insights into how people design long-horizon plans under ambiguity, with updating and learning to take place. Another direct extension is to analyse the data at the aggregate level. Instead of performing a subject level analysis, the data can be pooled so as to estimate a *mixture model* following the methodology that Conte and Hey (2013) applied. What

[16]Borrowing would create the same problem as in our experiment with the n0short selling constraint.

has been observed from the results of our experiments, is that while the theoretical models require remarkably high levels of sophistication from the decision makers, the subjects usually resort to simple heuristics in order to simplify the complexity of the tasks. This is a call for the creation of additional links between economic theory and psychology and a need for more *behavioural economics* and *bounded rationality* approaches to be developed. Psychologists have already developed similar models that remain inactive in the field of economics (see Busemeyer and Diederich (2002)). An interesting question that is raised is what is the cost in welfare of the various inconsistencies. At the end of the day, one should ask whether there is a cost of being dynamically inconsistent in terms of final wealth or altering the initial choices is a strategy that can maximise utility (the decision maker is better-off when is dynamically inconsistent). Finally, the updating issues can be extended to the cases where it is quite difficult to create priors. These cases are known as unforseen contingencies and appropriate models have been recently developed (see Karni and Viero (2013), Karni and Viero (2014)). It would be quite interesting to obtain insights of how people behave in similar environments. Nevertheless, it is difficult to apply this in the lab, as one should be very careful not to create suspicion or hints of deception in the experimental protocol to be applied.

The experiment on updating beliefs under ambiguity (Chapter 2) focused on individual choice, using a simple task to test several different updating rules and types of decision makers. Most of the applications in real economic life are being taken under the interaction or influence of other members in the society. Consequently, it is of paramount importance to extend the individual choice framework in order to capture similar interactions. The straightforward way to do this, is initially to extend the framework to a game-theoretical model where *strategy* now plays a significant role. Already in the game theoretical literature, the notion of Ellsberg games have been developed which aims to understand how the set of equilibria is extended or restricted when the players are characterised by ambiguity averse preferences (see Eichberger and Kelsey (2014)). Lately, there are extensions to dynamic games where the players decide either sequentially or multiple times, and at each decision mode (information set) there is some piece of information to allow updating. The standard way that the literature suggests to solve this is by assuming Savage players that are probabilistic sophisticated and model the game assuming a Bayesian-Nash equilibrium exists. Relaxing this assumption and allowing for non-Expected Utility preferences, for sure changes the set of equilibria (see Battigalli et al. (2013), Mouraviev et al. (2014) and Karni et al. (2013) for references). It is interesting to test which of those equilibria can behaviourally survive. Interaction is not constrained only on strategic and non-cooperative environments but also includes decisions that are made in the context of a group. There is very little experimental evidence of how groups form beliefs in the presence of ambiguity, how they update these beliefs and how they collectively decide. Most of the work has been done in risky environments. Keck et al. (2014) provide some initial results on this topic but further research is needed

On the methodological issues, there are three extensions that need to be done. On the experiment of chapter 3, the analysis is realised assuming that the decision variable follows a continuous distribution. As in the experiment the subjects are required to make the allocations in integer numbers, it is not unreasonable to assume that the distribution is actually discrete. Further modifications are needed in order to incorporate discrete beta distribution to the analysis and also to define the proposed specifications under a discrete distribution. Also, there are options to extend the specifications to more flexible and behaviourally oriented stories. A second methodological issue has to do with the elicitation of beliefs. In order to do so, we jointly estimated parameters that represent beliefs, preference parameters (risk aversion and weighting coefficients) and parameters of stochastic choice. It would be interesting to reduce the number of parameters by applying some different mechanism of beliefs' elicitation. In the literature, the most prominent is to use a *scoring rule* at the cost of not being able to assume non-linearities in the utility function. Kothiyal et al. (2010) present some possible extensions of how scoring rules can be extended to measure ambiguity and subjective beliefs. Empirical work is needed to validate the appropriateness of these methods. Finally, the experiment in chapter 3, provided evidence that subjects exhibit high degrees of risk aversion. What is not clear, is the issue of whether this was just a simple coincidence or allocation questions indeed make subjects to behave in a more risk averse way. Experimental tests would verify the correlation between the decision task and the degree of risk aversion elicitated.

Overall, if we would like to summarise the content of this thesis in a few lines, it would be that the standard economic model (SEU) is seriously challenged when decisions are made in a dynamic framework, people hold beliefs that are distorted due to ambiguity aversion, people significantly differ from each other regarding their preferences and their choices and further research is needed in order to understand how people form their prior beliefs, how do they update these priors in an individual and social context and what the consequences are in welfare by being dynamically inconsistent.

Bibliography

Abdellaoui, M. (2000), "Parameter-Free Elicitation of Utility and Probability Weighting Functions." *Management Science*, 46, pp. 1497–1512.

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

Ahn, D., S. Choi, D. Gale, and S. Kariv (2014), "Estimating Ambiguity Aversion in a Portfolio Choice Experiment." *Quantitative Economics*, 5, 195–223.

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

Andersen, S., J. Fountain, G. Harrison, and E. Rutstrom (2009), "Estimating Aversion to Uncertainty." Working Papers 07-2009, Copenhagen Business School, Department of Economics.

Andersen, S., J. Fountain, W. G. Harrison, A. Hole, and E. Rutstrom (2012), "Inferring Beliefs as Subjectively Imprecise Probabilities." *Theory and Decision*, 73, 161–184.

Antoniou, K., G. Harrison, I. Lau, and D. Read (2013), "Subjective Bayesian Beliefs." Working Papers wpn13-02, Warwick Business School, Finance Group.

Baillon, A., H Bleichrodt, U. Keskin, O. L'Haridon, and C. Li (2013), "Learning Under Ambiguity: An Experiment Using Initial Public Offerings on a Stock Market." Economics working paper archive, University of Rennes 1 & University of Caen.

Battigalli, Pierpaolo, Simone Cerreia-Vioglio, Fabio Maccheroni, and Massimo Marinacci (2013), "Mixed Extensions of Decision Problems under Uncertainty."

Working Papers 485, IGIER (Innocenzo Gasparini Institute for Economic Research), Bocconi University.

Becker, G., M. DeGroot, and J. Marschak (1964), "Measuring Utility by a Single-Response Sequential Method." *Behavioral Science*, 9(3), 226–32.

Bock, O., A. Nicklisch, and I. Baetge (2012), "hroot: Hamburg Registration and Organization Online Tool." Working Papers wpn13-02, H-Lab Working Paper No. 1, 2012.

Busemeyer, J., R. Barkan, X. Li, and Z. Ma (2000), "Dynamic and Consequential Consistency of Choices Between Paths of Decision Trees." *Journal of Experimental Psychology: General*, 129, 530–545.

Busemeyer, J. and A. Diederich (2002), "Survey of Decision Field Theory." *Mathematical Social Sciences*, 43, 345 – 370. Random Utility Theory and Probabilistic Measurement Theory.

Camerer, C. (1995), "Individual Decision Making." In *The Handbook of Experimental Economics* (J. Kagel and A. Roth, eds.), Princeton University Press.

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

Carbone, E. and G. Infante (2014), "Comparing Behavior Under Risk and Under Ambiguity in a Lifecycle Experiment." *Theory and Decision*, 77, 313–322.

Charness, G., E. Karni, and D. Levin (2007), "Individual and Group Decision Making under Risk: An Experimental Study of Bayesian Updating and Violations of First-Order Stochastic Dominance." *Journal of Risk and Uncertainty*, 35, 129–148.

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

Chateauneuf, A., T. Gajdos, and Jaffray J.Y. (2011), "Regular Updating." *Theory and Decision*, 71, 111–128.

Choi, S., R. Fisman, D. Gale, and S. Kariv (2007), “Consistency and heterogeneity of individual behavior under uncertainty.” *American Economic Review*, 97, 1921–1938.

Clarke, K. A. (2007), “A Simple Distribution-Free Test for Nonnested Model Selection.” *Political Analysis*, 15, 347–363.

Cohen, M., I. Gilboa, and D. Schmeidler (2000), “An Experimental Study of Updating Ambiguous Beliefs.” *Risk, Decision and Policy*, 5 (2), 123–133.

Conte, Anna and John Hey (2013), “Assessing Multiple Prior Models of Behaviour under Ambiguity.” *Journal of Risk and Uncertainty*, 46, 113–132.

Cubitt, P., R., C. Starmer, and R. Sugden (1998), “Dynamic Choice and the Common Ratio Effect: An Experimental Investigation.” *Economic Journal*, 108, 1362–80.

Cubitt, R. and R. Sugden (2001), “ Dynamic Decision-Making under Uncertainty: An Experimental Investigation of Choices between Accumulator Gambles.” *Journal of Risk and Uncertainty*, 22, 103–28.

Dempster, A. P. (1967), “Upper and Lower Probabilities Induced by a Multivalued Mapping.” *The Annals of Mathematical Statistics*, 38 (2), 325–339.

Dempster, A. P. (1968), “A Generalization of Bayesian Inference.” *Journal of the Royal Statistical Society. Series B (Methodological)*, 30, pp. 205–247.

Deryugina, T. (2013), “How Do People Update? The Effects of Local Weather Fluctuations on Beliefs about Global Warming.” *Climatic Change*, 118, 397–416.

Dominiak, A., P. Dürsch, and J. Lefort (2012), “A Dynamic Ellsberg Urn Experiment.” *Games and Economic Behavior*, 75, 625–638.

Eichberger, J., S. Grant, and D. Kelsey (2010), “Comparing Three Ways to Update Choquet Beliefs.” *Economics Letters*, 107, 91 – 94.

Eichberger, J., S. Grant, and D. Kesley (2007), “Updating Choquet Beliefs.” *Journal of Mathematical Economics*, 43, 888–899.

Eichberger, J. and D. Kelsey (2014), “Optimism and Pessimism in Games.” *International Economic Review*, 55, 483–505.

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

Epstein, L. and M. Le Breton (1993), "Dynamically Consistent Beliefs Must Be Bayesian." *Journal of Economic Theory*, 61, 1–22.

Epstein, L., J. Noor, and A. Sandroni (2010), "Non-Bayesian Learning." *The B.E. Journal of Theoretical Economics*, 10, 1–20.

Epstein, L. and M. Schneider (2003), "Recursive Multiple-Priors." *Journal of Economic Theory*, 113, 1–31.

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

Etner, J., M. Jeleva, and J.M. Tallon (2012), "Decision Theory Under Ambiguity." *Journal of Economic Surveys*, 26(2), 234–270.

Friedman, D. (1998), "Monty Hall's Three Doors: Construction and Deconstruction of a Choice Anomaly." *American Economic Review*, 88, 933–46.

Gajdos, T., T. Hayashi, J.-M. Tallon, and J.-C. Vergnaud (2008), "Attitude Toward Imprecise Information ." *Journal of Economic Theory*, 140, 27 – 65.

Ghalanos, Alexios and Stefan Theussl (2012), *Rsolnp: General Non-linear Optimization Using Augmented Lagrange Multiplier Method*. R package version 1.14.

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

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

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

Gilboa, I. (2009), *Theory of Decision under Uncertainty*. Cambridge University Press.

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

Gilboa, I. and D. Schmeidler (1993), "Updating Ambiguous Beliefs." *Journal of Economic Theory*, 59, 33–49.

Halevy, Y. (2007), "Ellsberg Revisited: An Experimental Study." *Econometrica*, 75, 503–536.

Hammond, J. and H. Zank (2014), "Rationality and Dynamic Consistency Under Risk and Uncertainty." In *Hanbook of the Economics of Risk and Uncertainty* (M. Machina and W. K. Viscusi, eds.), 41–97, Eslevier B.V.

Hammond, P. (1988), "Consequentialist Foundations for Expected Utility." *Theory and Decision*, 25, 25–78.

Hanany, E. and P. Klibanoff (2009), "Updating Ambiguity Averse Preferences." *The B.E. Journal of Theoretical Economics*, 9 (1), 1–53.

Hayashi, T. and R. Wada (2010), "Choice with Imprecise Information: an Experimental Approach." *Theory and Decision*, 69, 355–373.

Hey, J. (2014), "Choice Under Uncertainty: Empirical Methods and Experimental Results." In *Hanbook of the Economics of Risk and Uncertainty* (M. Machina and K. Viscusi, W., eds.), 809–850, Eslevier B.V.

Hey, J. and G. Lotito (2009), "Naive, Resolute or Sophisticated? A Study of Dynamic Decision Making." *Journal of Risk and Uncertainty*, 38, 1–25.

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

Hey, J. and N. Pace (2014), "The Explanatory and Predictive Power of Non Two-Stage-Probability Models of Decision Making Under Ambiguity." *Journal of Risk and Uncertainty*, Forthcoming.

Hey, J. and L. Panaccione (2011), "Dynamic Decision Making: What Do People Do?" *Journal of Risk and Uncertainty*, 42, 85–123.

Hey, J. and M.. Paradiso (1999), "Dynamic Choice and Timing-Independence: an Experimental Investigation." Discussion Papers 99/26, Department of Economics, University of York.

Hey, J. and M. Paradiso (2006), "Preferences Over Temporal Frames In Dynamic Decision Problems: An Experimental Investigation." *Manchester School*, 74, 123–137.

Holt, A., C. and M. Smith, A. (2009), "An Update on Bayesian Updating." *Journal of Economic Behavior & Organization*, 69, 125–134.

Holt, C. and S. Laury (2002), "Risk Aversion and Incentive Effects." *American Economic Review*, 92, 1644–1655.

Houser, D., M. Keane, and K. McCabe (2004), "Behavior in a Dynamic Decision Problem: An Analysis of Experimental Evidence Using a Bayesian Type Classification Algorithm." *Econometrica*, 72, 781–822.

Jaffray, J.Y. (2008), "Bayesian Updating and Belief Functions." In *Classic Works of the Dempster-Shafer Theory of Belief Functions* (RolandR. Yager and Liping Liu, eds.), volume 219 of *Studies in Fuzziness and Soft Computing*, 555–576, Springer Berlin Heidelberg.

Johnson, J. and R. Busemeyer, J. (2001), "Multiple-Stage Decision-Making: The Effect of Planning Horizon Length on Dynamic Consistency." *Theory and Decision*, 51, 217–246.

Karni, E., F. Macherroni, and M. Marinacci (2013), "Ambiguity and non-Expected Utility." Working papers, Working Paper.

Karni, E. and D. Schmeidler (1991), "Atemporal Dynamic Consistency and Expected Utility Theory." *Journal of Economic Theory*, 54, 401–408.

Karni, E. and M. Viero (2013), ""Reverse Bayesianism": A Choice-Based Theory of Growing Awareness." *American Economic Review*, 103, 2790–2810.

Karni, E. and M. Viero (2014), "Awareness of Unawareness: A Theory of Decision Making in the Face of Ignorance." Working papers, QED Working Paper No. 1322.

Keck, S., E. Diecidue, and D. Budescu (2014), "Group Decisions under Ambiguity: Convergence to Neutrality ." *Journal of Economic Behavior & Organization*, 103, 60 – 71.

Keynes, J. (1921), *A Treatise on Probability*. Macmillan And Co.

Klibanoff, P. and E. Hanany (2007), "Updating Preferences with Multiple Priors." *Theoretical Economics*, 2 (3), 261–298.

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

Klibanoff, P., M. Marinacci, and S. Mukerji (2009), "Recursive Smooth Ambiguity Preferences." *Journal of Economic Theory*, 144, 930–976.

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

Kothiyal, A., V. Spinu, and P. Wakker (2010), "Comonotonic Proper Scoring Rules to Measure Ambiguity and Subjective Beliefs." *Journal of Multi-Criteria Decision Analysis*, 17, 101–113.

Kothiyal, A., V. Spinu, and P. Wakker (2014), "An Experimental Test of Prospect Theory for Predicting Choice under Ambiguity." *Journal of Risk and Uncertainty*, 48, 1–17.

Loomes, G. (1991), "Evidence of a New Violation of the Independence Axiom." *Journal of Risk and Uncertainty*, 4, 91–108.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006a), "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences." *Econometrica*, 74, 1447–1498.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006b), "Dynamic Variational Preferences." *Journal of Economic Theory*, 128, 4–44.

Machina, M. (1989), "Dynamic Consistency and Non-expected Utility Models of Choice under Uncertainty." *Journal of Economic Literature*, 27, 1622–68.

Machina, M. (2013), "Ambiguity Aversion with Three or More Outcomes." Working papers, University of California.

Machina, M. and D. Schmeidler (1992), "A More Robust Definition of Subjective Probability." *Econometrica*, 60, 745–80.

McClennen, E.F. (1990), *Rationality and Dynamic Choice: Foundational Explorations*. Paperback re-issue, Cambridge University Press.

Moreno, O. and Y. Rosokha (2013), "Learning under Ambiguity an Experiment." Working papers, The University of Texas at Austin, Department of Economics.

Mouraviev, I., F. Riedel, and L. Sass (2014), "Kuhn's Theorem for Extensive Form Ellsberg Games." Working papers, Institute of Mathematical Economics Working Paper No. 510.

Mukerji, S. and J. Tallon (2004), "An Overview of Economic Applications of David Schmeidler's Models of Decision Making under Uncertainty." In *Uncertainty in Economic Theory: Essays in Honor of David Schmeidler's 65th Birthday* (I. Gilboa, ed.), Psychology Press.

Nebout, A. and M. Willinger (2014), "Are Non-Expected Utility Individual Really Dynamic Inconsistent? Experimental Evidence." Discussion Papers 99/26, LAMETA-CNRS, University of Montpellier.

Ozdenoren, E. and J. Peck (2008), "Ambiguity Aversion, Games Against Nature and Dynamic Consistency." *Games and Economic Behavior*, 62, 106–115.

Pires, C. (2002), "A Rule For Updating Ambiguous Beliefs." *Theory and Decision*, 53, 137–152.

Poinas, F., J. Rosaz, and B. Roussillon (2012), "Updating Beliefs with Imperfect Signals: Experimental Evidence." *Journal of Risk and Uncertainty*, 44, 219–241.

Pollak, R. (1968), "Consistent Planning." *Review of Economic Studies*, 35 (2), 201–208.

Prelec, D. (1998), "The Probability Weighting Function." *Econometrica*, 66, pp. 497–527.

Qiu, J. and U. Weitzel (2013), "Experimental Evidence on Valuation and Learning with Multiple Priors." Working papers, Radboud University Nijmegen.

Raiffa, H. (1968), *Decision Analysis – Introductory Lectures on Choices under Uncertainty*. Addison-Wesley, Reading, MA.

Savage, L. (1954), *The Foundations of Statistics*. Wiley, New York.

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

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

Selvin, S. (1975), "A Problem in Probability." *The American Statistician*, 29, 67–71.

Shafer, G. (1976), *A Mathematical Theory of Evidence*. Princeton University Press, Princeton.

Siniscalchi, M. (2009), "Vector Expected Utility and Attitudes Toward Variation." *Econometrica*, 77, 801–855.

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

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

Stott, H. (2006), "Cumulative Prospect Theory's Functional Menagerie." *Journal of Risk and Uncertainty*, 32, 101–130.

Strotz, R. (1955-56), "Myopia and Inconsistency in Dynamic Utility Maximization." *The Review of Economic Studies*, 23, 165–180.

Tversky, A. and D. Kahneman (1981), "The Framing of Decisions and the Psychology of Choice." *Science*, 211, 453–458.

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

Wakker, P. (1989), "Continuous Subjective Expected Utility with non-Additive Probabilities." *Journal of Mathematical Economics*, 18, 1–27.

Wakker, P. (1998), "Nonexpected Utility as Aversion of Information." *Journal of Behavioral Decision Making*, 1 (3), 169–175.

Wakker, P. (2007), "Message to Referees who Want to Embark on yet Another Discussion of the Random-Lottery Incentive System for Individual Choice." Working papers, University of Rotterdam.

Wakker, P. (2008), "Explaining the Characteristics of the Power (CRRA) Utility Family." *Health Economics*, 17, 1329–1344.

Wang, T. (2003), "Conditional Preferences and Updating." *Journal of Economic Theory*, 286–321.

Wilcox, N. (2008), "Stochastic Models for Binary Discrete Choice Under Risk: A Critical Stochastic ModeModel Primer and Econometric Comparison." In *Research in Experimental Economics* (J. Cox and G. Harrison, eds.), volume 12, 41–196, Emerald Group Publishing Limited.

Appendix A

Optimal Allocations for Chapter 2

This technical appendix contains the solutions for all the different preference functionals as well as the strategies that each type is applying. More specifically, there are the solutions for the Subjective Expected Utility (SEU) decision maker, the MaxMin Expected Utility (MEU), the Choquet Expected Utility (CEU) and the Source Choquet Expected Utility (SCEU). The solutions for the SEU, CEU and SCEU are analytical so the process is presented in full length while in the case of the MEU the optimal allocation is calculated with the help of numerical methods, so the specification of the problem, as well as the algorithm applied are presented.

A.1 Dynamic Consistency

This section provides the proof of the *strict* definition of dynamic consistency that we applied in chapter 2. This definition requires that the *ex-ante* choices are exactly the same with the *ex-post*.

Consider a two-stage decision problem. At the first stage the individual has m to allocate between three states of the world, with probabilities p_i, p_j and p_k which sum to one. Let the optimal allocations be x_i^*, x_j^* and x_k^*. Then the individual is told that state k has not occurred and thus loses x_k^*. The individual is then asked to allocate the remaining amount $x_i^* + x_j^*$ between states i and j. Let the optimal allocations be X_i^* and X_j^*. Then we want to show that $X_i^* = x_i^*$ and that $X_j^* = x_j^*$. This is a requirement of dynamic consistency, which is implied by EU.

The Lagrangian in the first problem is given by:

$$p_i u(x_i) + p_j u(x_j) + p_k u(x_k) + \lambda(m - x_i - x_j - x_k)$$

and hence the first-order conditions are

$$p_i u'(x_i^*) = \lambda$$
$$p_j u'(x_j^*) = \lambda$$
$$p_k u'(x_k^*) = \lambda$$
$$x_i^* + x_j^* + x_k^* = m$$

Note that the first two of these and the fourth imply that

$$p_i u'(x_i^*) = p_j u'(x_j^*) \tag{A.1}$$
$$x_i^* + x_j^* = m - x_k^* \tag{A.2}$$

Now the objective at the second stage is to choose X_i and X_j to maximise:

$$\frac{p_i}{p_i + p_j} u(X_i) + \frac{p_j}{p_i + p_j} u(X_j)$$

subject to $X_i + X_j = m - x_k^*$

This is the same (using a linear transformation) as choosing X_i and X_j to maximise:

$p_i u(X_i) + p_j u(X_j) + p_k u(x_k^*)$subject to $X_i + X_j + x_k^* = m$.

The Lagrangian is:

$p_i u(X_i) + p_j u(X_j) + p_k u(x_k^*) + \Lambda(m - X_i - X_j - x_k^*)$

and the first-order conditions are

$$p_i u'(X_i^*) = \Lambda$$
$$p_j u'(X_j^*) = \Lambda$$
$$X_i^* + X_j^* + x_k^* = m$$

These imply

$$p_i u'(X_i^*) = p_j u'(X_j^*) \tag{A.3}$$

$$X_i^* + X_j^* + x_k^* = m \tag{A.4}$$

Now note that equations A.1, A.2 are the same as equations A.3, A.4 and hence that $X_i^* = x_i^*$ and $X_j^* = x_j^*$. ■

A.2 Subjective Expected Utility (SEU)

In the case of SEU, there is no way to distinguish among resolute, naive and sophisticated decision makers as all types update (when updating takes place) using the Bayesian rule which guarantees *dynamic consistency*. Nevertheless, in this Appendix we provide the solution for all the three types by assuming SEU preferences. This is done for the simple reason that due to the elegance that the SEU model provides, concerning the algebra, it seems easier and more straightforward to describe the different strategies and the respective optimal allocations by assuming Expected Utility maximisation. Then, this methodology can be extended and applied to all the non-Expected Utility representations that we consider.

A.2.1 Resolute

A *resolute* decision maker, solves the problem as if it is a one-period problem. The allocation that is made at stage 1 based on the subjective probabilities, coincides with the conditional allocations. The problem of a resolute decision maker can be solved in two ways, which both lead to the same result. The first is to consider only the first stage allocation and solve for the optimal levels of x_i^*, x_j^*, x_k^* given the budget constraint. The objective function to maximise is:

$$\max_{x_i^*, x_j^*, x_k^*} p_i u(e_i x_i) + p_j u(e_j x_j) + p_k u(e_k x_k)$$

$$\text{s.t. } x_i^* + x_j^* + x_k^* = m$$

Assuming a Relative Absolute Risk Aversion (CRRA) function of the following form:

$$u(x) = \begin{cases} x^r & \text{if } r > 0 \\ \ln(x) & \text{if } r = 0 \\ -x^r & \text{if } r < 0 \end{cases}$$

The Lagrangian is:

$$\mathcal{L} = p_i(e_i x_i)^r + p_j(e_j x_j)^r + p_k(e_k x_k)^r + \lambda(m - x_i - x_j - x_k)$$

deriving the first order conditions:

$$\mathcal{L}_{x_i} = 0 \Rightarrow p_i e_i (e_i x_i)^{1-r} = \lambda$$

$$\mathcal{L}_{x_j} = 0 \Rightarrow p_j e_j (e_j x_j)^{1-r} = \lambda$$

$$\mathcal{L}_{x_k} = 0 \Rightarrow p_k e_k (e_k x_k)^{1-r} = \lambda$$

and using the budget constraint we obtain the optimal unconditional allocations:

$$x_i^* = \frac{m(p_i e_i)^{1/(1-r)} e_j e_k}{(p_i e_i)^{1/(1-r)} e_j e_k + (p_j e_j)^{1/(1-r)} e_i e_k + (p_k e_k)^{1/(1-r)} e_i e_j}$$

$$x_j^* = \frac{m(p_j e_j)^{1/(1-r)} e_i e_k}{(p_i e_i)^{1/(1-r)} e_j e_k + (p_j e_j)^{1/(1-r)} e_i e_k + (p_k e_k)^{1/(1-r)} e_i e_j}$$

$$x_k^* = \frac{m(p_k e_k)^{1/(1-r)} e_i e_j}{(p_i e_i)^{1/(1-r)} e_j e_k + (p_j e_j)^{1/(1-r)} e_i e_k + (p_k e_k)^{1/(1-r)} e_i e_j}$$

An alternative way to solve this problem, is to think in a *backward induction*[1] way and to consider all the possible future states of the world. Stating this in a different way, to distinguish among the different conditional allocations (for instance, the allocation to j conditional that the state is not i is different to the allocation to j conditional to the information that the state is not k. This leads to six different conditional allocations (two for each not-state) which are denoted as $x_{j\neg i}, x_{k\neg i}, x_{i\neg j}, x_{k\neg j}, x_{i\neg k}, x_{j\neg k}$ where $x_{j\neg i}$ stands for the allocation to asset j when the information is that the state of the world is not i . The maximisation problem to be solved now is:

[1]This method should not be confused with the sophisticated type that is presented later. The idea in the method, is that the decision maker anticipated that there will be conditional states, nevertheless, she is solving the problem as a static one considering all the conditional states.

$$\max_{\{x_{i\neg j}, x_{i\neg k}, x_{j\neg i}, x_{j\neg k}, x_{k\neg i}, x_{k\neg j}\}} (q_{j\neg i})u(e_j x_{j\neg i}) + (q_{k\neg i})u(e_k x_{k\neg i}) + (q_{i\neg j})u(e_i x_{i\neg j}) + \tag{A.5}$$

$$(q_{k\neg j})u(e_k x_{k\neg j}) + (q_{i\neg k})u(e_i x_{i\neg k}) + (q_{j\neg k})u(e_j x_{j\neg k}) \tag{A.6}$$

$$\text{s.t. } x_{i\neg j} + x_{i\neg k} + x_{j\neg i} + x_{j\neg k} + x_{k\neg i} + x_{k\neg j} = 2m \tag{A.7}$$

where now $q_{j\neg i}$ stands for the compound probability which is simply the conditional probability of j being the real state of the world when the event *not i* has occurred multiplied by the probability that the event *not i* may happen. In this example, this compound probability can be written as:

$$q_{j\neg i} = P(s_j|\neg i)P(\neg i)\frac{P(s_j)}{P(s_j) + P(s_k)}(P(s_j) + P(s_k))$$

which is equal to $P(s_j)$.[2]

The first order conditions are:

$$\mathcal{L}_{x_{i\neg j}} = 0 \Rightarrow q_{i\neg j} e_i (e_i x_{i\neg j})^{-1/r} = \lambda_1 \tag{A.8}$$

$$\mathcal{L}_{x_{i\neg k}} = 0 \Rightarrow q_{i\neg k} e_i (e_i x_{i\neg k})^{-1/r} = \lambda_1 \tag{A.9}$$

$$\mathcal{L}_{x_{j\neg i}} = 0 \Rightarrow q_{j\neg i} e_j (e_j x_{j\neg i})^{-1/r} = \lambda_1 \tag{A.10}$$

$$\mathcal{L}_{x_{j\neg k}} = 0 \Rightarrow q_{j\neg k} e_j (e_j x_{j\neg k})^{-1/r} = \lambda_1 \tag{A.11}$$

$$\mathcal{L}_{x_{k\neg i}} = 0 \Rightarrow q_{k\neg i} e_k (e_k x_{k\neg i})^{-1/r} = \lambda_1 \tag{A.12}$$

$$\mathcal{L}_{x_{k\neg j}} = 0 \Rightarrow q_{k\neg j} e_k (e_k x_{k\neg j})^{-1/r} = \lambda_1 \tag{A.13}$$

Using the budget constraint we can solve for the conditional allocations:

[2] This result will be later useful when we will extend the analysis to the updating of non-additive capacities where the compound and the unconditional probabilities do not necessarily coincide.

$$x_{i\neg j} = \frac{2m(q_{i\neg j}e_i)^r e_j e_k}{(q_{j\neg i}e_j)^r e_i e_k + (q_{k\neg i}e_k)^r e_j e_i + (q_{i\neg j}e_i)^r e_j e_k + (q_{k\neg j}e_k)^r e_i e_j + (q_{i\neg k}e_i)^r e_j e_k + (q_{j\neg k}e_j)^r e_i e_k} \tag{A.14}$$

$$x_{i\neg k} = \frac{2m(q_{i\neg k}e_i)^r e_j e_k}{(q_{j\neg i}e_j)^r e_i e_k + (q_{k\neg i}e_k)^r e_j e_i + (q_{i\neg j}e_i)^r e_j e_k + (q_{k\neg j}e_k)^r e_i e_j + (q_{i\neg k}e_i)^r e_j e_k + (q_{j\neg k}e_j)^r e_i e_k} \tag{A.15}$$

$$x_{j\neg i} = \frac{2m(q_{j\neg i}e_j)^r e_i e_k}{(q_{j\neg i}e_j)^r e_i e_k + (q_{k\neg i}e_k)^r e_j e_i + (q_{i\neg j}e_i)^r e_j e_k + (q_{k\neg j}e_k)^r e_i e_j + (q_{i\neg k}e_i)^r e_j e_k + (q_{j\neg k}e_j)^r e_i e_k} \tag{A.16}$$

$$x_{j\neg k} = \frac{2m(q_{j\neg k}e_j)^r e_i e_k}{(q_{j\neg i}e_j)^r e_i e_k + (q_{k\neg i}e_k)^r e_j e_i + (q_{i\neg j}e_i)^r e_j e_k + (q_{k\neg j}e_k)^r e_i e_j + (q_{i\neg k}e_i)^r e_j e_k + (q_{j\neg k}e_j)^r e_i e_k} \tag{A.17}$$

$$x_{k\neg i} = \frac{2m(q_{k\neg i}e_k)^r e_i e_j}{(q_{j\neg i}e_j)^r e_i e_k + (q_{k\neg i}e_k)^r e_j e_i + (q_{i\neg j}e_i)^r e_j e_k + (q_{k\neg j}e_k)^r e_i e_j + (q_{i\neg k}e_i)^r e_j e_k + (q_{j\neg k}e_j)^r e_i e_k} \tag{A.18}$$

$$x_{k\neg j} = \frac{2m(q_{k\neg j}e_k)^r e_i e_j}{(q_{j\neg i}e_j)^r e_i e_k + (q_{k\neg i}e_k)^r e_j e_i + (q_{i\neg j}e_i)^r e_j e_k + (q_{k\neg j}e_k)^r e_i e_j + (q_{i\neg k}e_i)^r e_j e_k + (q_{j\neg k}e_j)^r e_i e_k} \tag{A.19}$$

$$\tag{A.20}$$

It is easy to notice that $x_{i\neg j} = x_{i\neg k}$ as the updating is realised according to the Bayesian rule, which leads to the same allocation as the two formulas are equivalent. This type of decision maker is introduced in the literature by Machina (1989) and later by McClennen (1990) where the main idea behind this, is that of commitment. Also this kind of behaviour seems to be consistent with actions that are contrary to *consequentialism* and that in fact take into consideration the history of actions. A decision maker decides to impose the decisions that she made at the first stage and thus, she is not willing to make any changes at the second stage. Thus, the decision that is made in the first stage is the one that is implemented throughout the following steps. The decision maker does not rearrange her initial decisions and does not use the additional information that is revealed (Al Najjar (2009) refers to this as "ignoring relevant information").

A.2.2 Sophisticated Type

This type of decision maker solves the problem in two steps using *backward induction*. Solving backward, in the first step, the decision maker thinks what she would do if she was to reach a specific decision node. It is useful to use again the decision tree that represents the problem. In this problem there are three

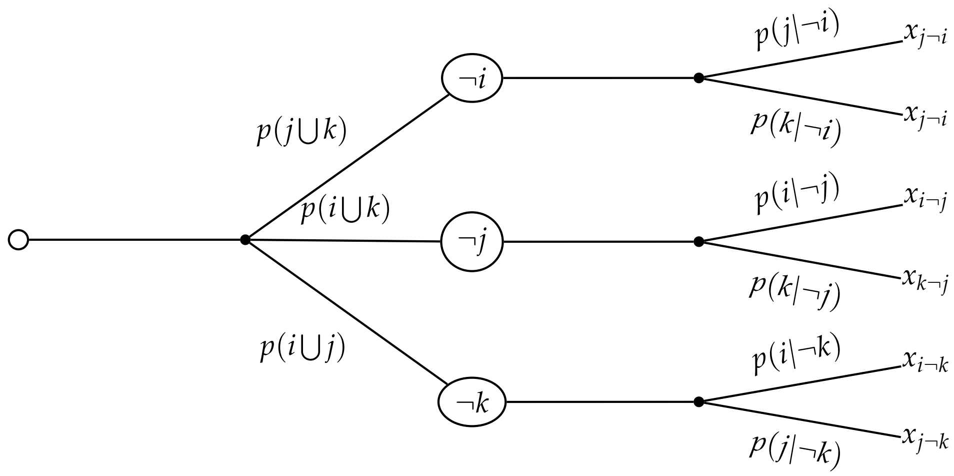

Figure A.1 The 2-Stage Decision Tree

decision nodes (*not colour s*) which form the set $\{s_{\neg i}, s_{\neg j}, s_{\neg k}\}$. At each decision node, the objective is to solve for the optimal conditional allocation taking as given the available income (which is the sum of the optimal amounts that were allocated at stage 1 of the experiment $(x_j^* + x_k^*)$ if the state of the world is $s_{\neg i}$). This leads to the solution for the conditional allocations as a function of the available *conditional* income[3] In the second step, taking into consideration these conditional allocations, the decision maker solves for the unconditional optimal allocation x_i^*, x_j^*, x_k^* which is exactly what she does at the initial stage. In this experiment, the decision maker is liable to face three different events. Therefore, she has to solve three different conditional maximisation problems in the second stage, and the solutions from these problems will be used to define the initial allocations at stage 1. This leads to six different conditional allocations as these are all the possible states of the world that may be realised. We denote $x_{j\neg i}$ the conditional allocation to j and $p(s_j|s_{\neg i})$ the conditional probability of state j, when the information *the ball is*

[3]Here there are two ways that we can proceed. The first is by approaching the solution using the conditional income $m_{\neg i} = x_j{}^* + x_k{}^*$ as the variable that we are solving for, or by using the General Envelope theorem, we can solve for the optimal allocations $x_i{}^*$. We provide the solution for both.

not i is provided. Consequently, there are six conditional allocations denoted as $x_{j\neg i}, x_{k\neg i}, x_{i\neg j}, x_{k\neg j}, x_{i\neg k}, x_{j\neg k}$ where $x_{j\neg i}$ is the amount that is allocated to the asset j when the state is *not i*. .

We can now provide the analytical solutions for the amounts allocated at stage 1, $\{x_i^*, x_j^*, x_k^*\}$ and the six conditional allocations $\{x_{j\neg i}, x_{k\neg i}, x_{i\neg j}, x_{k\neg j}, x_{i\neg k}, x_{j\neg k}$ where $x_{j\neg i}\}$.

Notice that in the case of subjective Expected Utility, as we will show later, $x_{i\neg j}$ and $x_{i\neg k}$ will lead to exactly the same solution as x_i^* due to the application of the Bayesian rule.

Consider the case where in the second stage, the information revealed is *the ball is not i* thus $s_{\neg i}$, the problem in the general form can be written as:

$$\max_{\{x_{j\neg i}, x_{k\neg i}\}} p(s_j|s_{\neg i})u(e_j x_{j\neg i}) + p(s_k|s_{\neg i})u(e_k x_{k\neg i})$$

$$\text{s.t. } x_{j\neg i} + x_{k\neg i} = x_j^* + x_k^*$$

Despite the fact that in the case of Expected Utility the conditional probabilities are much simplified, it is useful to keep the conditional notation, as in the case of non-additive capacities this is not true any more. At the first stage, the decision maker solves the following three problems:
When the ball is not i:

$$\max_{x_{j\neg i}, x_{k\neg i}} p(s_j|\neg i)u(z_{j\neg i}) + p(s_k|\neg i)u(z_{k\neg i}) \tag{A.21}$$

$$\text{s.t. } e_j x_{j\neg i} - e_k x_{k\neg i} \leq 0 \tag{A.22}$$

$$x_{j\neg i} + x_{k\neg i} = x_j^* + x_k^* \tag{A.23}$$

where $z_{j\neg i} = e_j x_{j\neg i}$ and $e_k x_{k\neg i}$ the payoffs at each case. We also write the conditional income as $m_{\neg i} = x_j^* + x_k^*$. Also for simplicity, the conditional probability $p(s_j|\neg i)$ will be denoted as $pj\neg i$. The two remaining conditional problems are symmetric to the one presented above. The solutions of these maximisation problems give the optimal conditional allocations as a function of the optimal unconditional allocations $x_{j\neg i} = x_{j\neg i}(x_j^* + x_k^*)$. Let $m_{\neg i} = x_j^* + x_k^*$ be the *conditional income* for each of the states. The problem can now be solved in two ways, as for $m_{\neg i}$ or as for $x_j^* + x_k^*$. We adopt the former. Forming the respective Lagrangian functions and solving for the conditional allocations of the first step we obtain:

$$x^*_{j\neg i} = \frac{m_{\neg i} e_k (e_k w'_k)^{1/(r-1)}}{e_j (e_j w'_j)^{1/(r-1)} + e_k (e_k w'_k)^{1/(r-1)}} \tag{A.24}$$

and similarly the conditional allocation to k:

$$x^*_{k\neg i} = \frac{m_{\neg i} e_j (e_j w'_j)^{1/(r-1)}}{e_j (e_j w'_j)^{1/(r-1)} + e_k (e_k w'_k)^{1/(r-1)}} \tag{A.25}$$

or simply $x^*_{k\neg i} = m_{\neg i} - x^*_{j\neg i}$. w'_s stands for the respective conditional probability.

Following the same procedure for the other two conditional states we obtain the conditional allocations. The solution is omitted as it is symmetric to the previous one.

These are the optimal conditional allocations which are functions of the conditional income $m_{\neg s}$ at each state. In the second step, the decision maker solves for the optimal levels of the three conditional incomes $m_{\neg i}, m_{\neg j}, m_{\neg k}$, taking as given the conditional allocations from step 1. Before writing the objective function, recall that $p_{\neg i}$ stands for the unconditional probability that the state of the world is not i and is equal to $p_j + p_k$. With $q_{j\neg i}$ we denote the compound probability that state j will happen when the state of the world is not i, which is simple the product between the unconditional probability that the state will be either j or k and the conditional probability of getting j when the event is that the state is not i. The formula for this compound probability is $q_{j\neg i} = p_{\neg i} p_{j\neg i} \Rightarrow q_{j\neg i} = p_{\neg i} \frac{p_j}{p_{\neg i}}$. The objective function can now be written as:

$$\begin{aligned}\mathcal{L} = \frac{1}{2}[q_{j\neg i} u(e_j x_{j\neg i}(m_{\neg i})) + q_{k\neg i} u(e_k x_{k\neg i}(m_{\neg i})) + q_{i\neg j} u(e_i x_{i\neg j}(m_{\neg j})) \\ + q_{k\neg j} u(e_k x_{k\neg j}(m_{\neg j})) + q_{i\neg k} u(e_i x_{i\neg k}(m_{\neg k})) + q_{j\neg k} u(e_j x_{j\neg k}(m_{\neg k}))] \\ + \lambda(2m - m_{\neg i}{}^* - m_{\neg j}{}^* - m_{\neg k}{}^*)\end{aligned} \tag{A.26}$$

and substituting with the optimal conditional allocations and the appropriate form of the utility function:

$$\mathcal{L} = \frac{1}{2}[q_{j\neg i}(e_j x^*_{j\neg i}(m_{\neg i}))^r + q_{k\neg i}(e_k x^*_{k\neg i}(m_{\neg i}))^r + q_{i\neg j}(e_i x^*_{i\neg j}(m_{\neg j}))^r \quad (A.27)$$
$$+q_{k\neg j}(e_k x^*_{k\neg j}(m_{\neg j}))^r + q_{i\neg k}(e_i x^*_{i\neg k}(m_{\neg k}))^r + q_{j\neg k}(e_j x^*_{j\neg k}(m_{\neg k}))^r]$$
$$+\lambda(2m - m_{\neg i}{}^* - m_{\neg j}{}^* - m_{\neg k}{}^*)$$

Deriving the FOC[4]:

$$\mathcal{L}_{m_{\neg i}{}^*} = 0 \Rightarrow m^{r-1}_{\neg i} \frac{e_j e_k}{e_j(e_j p_{j\neg i})^{\frac{1}{1-r}} + e_k(e_k p_{k\neg i})^{\frac{1}{1-r}}} (q_{j\neg i}(e_k p_{k\neg i})^r + q_{k\neg i}(e_j p_{j\neg i})^r) = \mu \quad (A.28)$$

$$\mathcal{L}_{m_{\neg j}{}^*} = 0 \Rightarrow m^{r-1}_{\neg j} \frac{e_i e_k}{e_i(e_i p_{i\neg j})^{\frac{1}{1-r}} + e_k(e_k p_{k\neg j})^{\frac{1}{1-r}}} (q_{i\neg j}(e_k p_{k\neg j})^r + q_{k\neg j}(e_i p_{i\neg j})^r) = \mu \quad (A.29)$$

$$\mathcal{L}_{m_{\neg k}{}^*} = 0 \Rightarrow m^{r-1}_{\neg k} \frac{e_i e_j}{e_j(e_j p_{j\neg k})^{\frac{1}{1-r}} + e_i(e_i p_{i\neg k})^{\frac{1}{1-r}}} (q_{j\neg k}(e_i p_{i\neg k})^r + q_{i\neg k}(e_j p_{j\neg k})^r) = \mu \quad (A.30)$$

By setting

$$A = [\frac{e_j e_k}{e_j(e_j p_{j\neg i})^{\frac{1}{1-r}} + e_k(e_k p_{k\neg i})^{\frac{1}{1-r}}} (q_{j\neg i}(e_k p_{k\neg i})^r + q_{k\neg i}(e_j p_{j\neg i})^r)]^{\frac{1}{r-1}}$$

$$B = [\frac{e_i e_k}{e_i(e_i p_{i\neg j})^{\frac{1}{1-r}} + e_k(e_k p_{k\neg j})^{\frac{1}{1-r}}} (q_{i\neg j}(e_k p_{k\neg j})^r + q_{k\neg j}(e_i p_{i\neg j})^r)]^{\frac{1}{r-1}}$$

$$C = [\frac{e_i e_j}{e_j(e_j p_{j\neg k})^{\frac{1}{1-r}} + e_i(e_i p_{i\neg k})^{\frac{1}{1-r}}} (q_{j\neg k}(e_i p_{i\neg k})^r + q_{i\neg k}(e_j p_{j\neg k})^r)]^{\frac{1}{r-1}}$$

equating A.28 to A.29 and A.28 to A.30 and using the constraint that $(m_{\neg i} + m_{\neg j} + m_{\neg k} = 2m)$, we can solve for the optimal conditional incomes as a function of the initial income m, the exchange rates e_i,e_j,e_k the coefficient of risk aversion r and

[4]Here the problem can be solved in two ways. The first is to define $m_{\neg i} = x_j{}^* + x_k{}^*$ and take the derivative with respect to each conditional income. This will lead to 3 equations plus the budget constraint from which we can solve for the individual allocations. The second methodology is to use the *General Envelop Theorem* and to prove that the derivatives from the first methodology are equivalent to the partial derivatives with respect to each optimal allocation $x_i{}^*$. Using either approach will lead to the same first order conditions and consequently to the same optimal allocations.

the conditional probabilities.

$$m_{\neg i} = \frac{2m}{1 + \frac{A}{B} + \frac{A}{C}}$$

$$m_{\neg j} = \frac{2m}{1 + \frac{B}{A} + \frac{B}{C}}$$

$$m_{\neg k} = \frac{2m}{1 + \frac{C}{A} + \frac{C}{B}}$$

The optimal allocations can then be simply found by using the budget constraint. This equation writes:

$$x_i^* + x_j^* + x_k^* = m$$

and as we defined earlier the conditional incomes, it is easy to solve for the optimal allocations. This gives the following equations:

$$x_i^* = m - m_{\neg i}$$

$$x_j^* = m - m_{\neg j}$$

$$x_k^* = m - m_{\neg k}$$

Then, we can use the solutions for the optimal conditional incomes and substitute back to the equations of the conditional allocations that we derived in the first stage of the maximisation problem. As the update is done using the Bayesian rule, this ensures Dynamic Consistency and thus the conditional allocations (*ex-post*) will be identical to the initial ones (*ex-ante*).

A.3 MaxMin Expected Utility (MEU)

For the multiple priors model we used numerical optimisation techniques to obtain the optimal solutions. It is a computationally intense task that requires several calculations in order to converge to the optimal allocation. For the *resolute* type, notice that as we are not aware where exactly the priors lie inside the triangle, we need to find the vector $x = \{x_i^*, x_j^*, x_k^*\}$ that maximises the minimum level of Expected Utility. To do so, based on the triangle illustrated by Figure 2.2 we apply this calculation using the prior sets that are formed at each of the vertices of the small triangle. The maximisation problem can be written as:

$$\max_{\{x_i^*, x_j^*, x_k^*\}} (\min(EU(P_A, x), EU(P_B, x), EU(P_C, x))) \tag{A.31}$$

and substituting the appropriate priors:

$$\begin{aligned}\max_{\{x_i^*, x_j^*, x_k^*\}} (\min((&\underline{p_i}(e_k x_i)^r + \underline{p_j}(e_j x_j)^r + (1-\underline{p_i}-\underline{p_j})(e_k(m-x_i-x_j))^r),\\ &(\underline{p_i}(e_k x_i)^r + (1-\underline{p_i}-\underline{p_k})(e_j x_j)^r + \underline{p_k}(e_k(m-x_i-x_j))^r),\\ &((1-\underline{p_j}-\underline{p_k})(e_k x_i)^r + \underline{p_j}(e_j x_j)^r + \underline{p_k}(e_k(m-x_i-x_j))^r)))\end{aligned}$$

where the budget constraint is directly substituted in the objective function and now it suffices to solve only for the optimal levels of i and j.

A *naive* decision maker, firstly solves A.31. Then receives the partial information (e.g. not j) and solves the second stage problem, based on the conditional income acquired at stage 1 and the updated beliefs. At stage 2, she solves:

$$\begin{aligned}\max_{\{x_{i\neg j}, x_{k\neg j}\}} &\{\min\{w_i(u(x_{i\neg j}) + (1-w_i)u(x_{k\neg j}), w_k(u(x_{i\neg j}) + (1-w_k)u(x_{k\neg j})\}\}\\ &\text{s.t. } x_{i\neg j} + x_{k\neg j} = x_i^* + x_k^*\end{aligned} \tag{A.32}$$

The same process is also applied to the other two conditional states. In the case of the conditional updating, as was shown in the main text, we do not need to consider the prior that maximises the likelihood of the foregone state. Consequently, we consider only two priors and search where the minimum utility is maximised there. For a *sophisticated* decision maker, the problem at the first step requires to solve for the optimal allocation given that some information for the possible state of the world has been received. The three conditional problems of states i, j and k are respectively:

$$\max(\min((1-\underline{p_i}-\underline{p_k})/(1-\underline{p_i})(e_j x_{j\neg i})^r + \underline{p_k}/(1-\underline{p_i})(e_k(m-x_i-x_{j\neg i}))^r,$$
$$\underline{p_j}/(1-\underline{p_i})(e_j x_{j\neg i})^r + (1-\underline{p_i}-\underline{p_j})/(1-\underline{p_i})(e_k(m-x_i-x_{j\neg i}))^r)) \quad \text{(A.33)}$$

$$\max(\min((1-\underline{p_j}-\underline{p_k})/(1-\underline{p_j})(e_i x_{i\neg j})^r + \underline{p_k}/(1-\underline{p_j})(e_k(m-x_j-x_{i\neg j}))^r,$$
$$\underline{p_i}/(1-\underline{p_j})(e_i x_{i\neg j})^r + (1-\underline{p_i}-\underline{p_j})/(1-\underline{p_j})(e_k(m-x_j-x_{i\neg j}))^r)) \quad \text{(A.34)}$$

$$\max(\min((1-\underline{p_j}-\underline{p_k})/(1-\underline{p_k})(e_i x_{i\neg k})^r + \underline{p_j}/(1-\underline{p_k})(e_j(m-x_k-x_{i\neg k}))^r,$$
$$\underline{p_i}/(1-\underline{p_k})(e_i x_{i\neg k})^r + (1-\underline{p_i}-\underline{p_k})/(1-\underline{p_k})(e_j(m-x_k-x_{i\neg k}))^r)) \quad \text{(A.35)}$$

and then solves for the problem as seen from period 1 to find the optimal levels for x_i^*, x_j^*, taking into consideration the optimal solutions of the conditional problems at period 2, given by A.33, A.34 and A.35. The problem now to solve is:

$$\max(\min(.5((\underline{p_k}+\underline{p_j})(\min((1-\underline{p_i}-\underline{p_k})/(1-\underline{p_i})(e_jx_{j\neg i})^r+\underline{p_k}/(1-\underline{p_i})(e_k(x_{k\neg i}))^r,$$
$$\underline{p_j}/(1-\underline{p_i})(e_jx_{j\neg i})^r+(1-\underline{p_i}-\underline{p_j})/(1-\underline{p_i})(e_k(x_{k\neg i}))^r))$$
$$+(1-\underline{p_j})(\min((1-\underline{p_j}-\underline{p_k})/(1-\underline{p_j})(e_ix_{i\neg j})^r+\underline{p_k}/(1-\underline{p_j})(e_k(x_{k\neg j}))^r,$$
$$\underline{p_i}/(1-\underline{p_j})(e_ix_{i\neg j})^r+(1-\underline{p_i}-\underline{p_j})/(1-\underline{p_j})(e_k(x_{k\neg j}))^r))$$
$$+(1-\underline{p_k})(\min((1-\underline{p_j}-\underline{p_k})/(1-\underline{p_k})(e_ix_{i\neg k})^r+\underline{p_j}/(1-\underline{p_k})(e_j(x_{j\neg k}))^r,$$
$$\underline{p_i}/(1-\underline{p_k})(e_ix_{i\neg k})^r+(1-\underline{p_i}-\underline{p_k})/(1-\underline{p_k})(e_j(x_{j\neg k}))^r))),$$

$$.5((1-\underline{p_i})(\min((1-\underline{p_i}-\underline{p_k})/(1-\underline{p_i})(e_jx_{j\neg i})^r+\underline{p_k}/(1-\underline{p_i})(e_k(m-x_i-x_{j\neg i}))^r,$$
$$\underline{p_j}/(1-\underline{p_i})(e_jx_{j\neg i})^r+(1-\underline{p_i}-\underline{p_j})/(1-\underline{p_i})(e_k(m-x_i-x_{j\neg i}))^r))$$
$$+(\underline{p_i}+\underline{p_k})(\min((1-\underline{p_j}-\underline{p_k})/(1-\underline{p_j})(e_ix_{i\neg j})^r+\underline{p_k}/(1-\underline{p_j})(e_k(x_{k\neg j}))^r,$$
$$\underline{p_i}/(1-\underline{p_j})(e_ix_{i\neg j})^r+(1-\underline{p_i}-\underline{p_j})/(1-\underline{p_j})(e_k(x_{k\neg j}))^r))$$
$$+(1-\underline{p_k})(\min((1-\underline{p_j}-\underline{p_k})/(1-\underline{p_k})(e_ix_{i\neg k})^r+\underline{p_j}/(1-\underline{p_k})(e_j(x_{j\neg k}))^r,$$
$$\underline{p_i}/(1-\underline{p_k})(e_ix_{i\neg k})^r+(1-\underline{p_i}-\underline{p_k})/(1-\underline{p_k})(e_j(x_{j\neg k}))^r))),$$

$$.5((1-\underline{p_i})(\min((1-\underline{p_i}-\underline{p_k})/(1-\underline{p_i})(e_jx_{j\neg i})^r+\underline{p_k}/(1-\underline{p_i})(e_k(m-x_i-x_{j\neg i}))^r,$$
$$\underline{p_j}/(1-\underline{p_i})(e_jx_{j\neg i})^r+(1-\underline{p_i}-\underline{p_j})/(1-\underline{p_i})(e_k(m-x_i-x_{j\neg i}))^r))$$
$$+(1-\underline{p_j})(\min((1-\underline{p_j}-\underline{p_k})/(1-\underline{p_j})(e_ix_{i\neg j})^r+\underline{p_k}/(1-\underline{p_j})(e_k(x_{k\neg j}))^r,$$
$$\underline{p_i}/(1-\underline{p_j})(e_ix_{i\neg j})^r+(1-\underline{p_i}-\underline{p_j})/(1-\underline{p_j})(e_k(x_{k\neg j}))^r))$$
$$+(\underline{p_i}+\underline{p_j})(\min((1-\underline{p_j}-\underline{p_k})/(1-\underline{p_k})(e_ix_{i\neg k})^r+\underline{p_j}/(1-\underline{p_k})(e_j(x_{j\neg k}))^r,$$
$$\underline{p_i}/(1-\underline{p_k})(e_ix_{i\neg k})^r+(1-\underline{p_i}-\underline{p_k})/(1-\underline{p_k})(e_j(x_{j\neg k}))^r)))))$$

A.4 Choquet Expected Utility (CEU)

The problem to solve in its general form is to maximise the Choquet Expected Utility subject to the budget constraint and to the different rankings of the outcomes. This means that each ranking, defines a different constrained optimisation problem which requires the use of the Kuhn-Tucker conditions. We start with the case of a resolute decision maker that solves the problem as a static one and at the second stage applies the same allocations as she did at the first stage. In the case of CEU, as the beliefs are non-additive, we need to take into consideration the individual capacities for each state of the world as well as the joint capacities. In total there are six different capacities to estimate, namely v_i, v_j, v_k the individual capacities and v_{ij}, v_{ik}, v_{jk} the joint capacities.

In the case of a *resolute* decision maker, there is no updating process. For a resolute decision maker there are 13 different ranking of the outcomes. Solving for these 13 different rankings we obtain the equations for the optimal allocations. We remind that for the ranking $z_k > z_j > z_i$, where $z_s = e_s x_s$, the weights are defined in the following way:

$$\begin{aligned} \omega_k &= v_k \\ \omega_j &= v_{jk} - v_k \\ \omega_i &= 1 - v_{jk} \end{aligned}$$

Table A.1 summarises the weights for each of the possible rankings.

Table A.1 Table with all possible rankings

	Ranking	ω_i	ω_j	ω_k
1	$z_i > z_j > z_k$	v_i	$v_{ij} - v_i$	$1 - v_{ij}$
2	$z_i > z_k > z_j$	v_i	$1 - v_{ik}$	$v_{ik} - v_i$
3	$z_j > z_i > z_k$	$v_{ij} - v_j$	v_j	$1 - v_{ij}$
4	$z_j > z_k > z_i$	$1 - v_{jk}$	v_j	$v_{jk} - v_j$
5	$z_k > z_i > z_j$	$v_{ik} - v_k$	$1 - v_{ik}$	v_k
6	$z_k > z_j > z_i$	$1 - v_{jk}$	$v_{jk} - v_k$	v_k
7	$z_i = z_j > z_k$	v_i	$v_{ij} - v_i$	$1 - v_{ij}$
8	$z_i = z_k > z_j$	v_i	$1 - v_{ik}$	$v_{ik} - v_i$
9	$z_j = z_k > z_i$	$1 - v_{jk}$	v_j	$v_{jk} - v_j$
10	$z_i > z_j = z_k$	v_i	$v_{ij} - v_i$	$1 - v_{ij}$
11	$z_j > z_i = z_k$	$v_{ij} - v_j$	v_j	$1 - v_{ij}$
12	$z_k > z_i = z_j$	$v_{ik} - v_k$	$1 - v_{ik}$	v_k
13	$z_i = z_j = z_k$	$1/3$	$1/3$	$1/3$

The objective is to maximise:

$$\max_{x_i^*, x_j^*, x_k^*} \omega_i u(z_i) + \omega_j u(z_j) + \omega_k u(z_k)$$

subject to the budget constraint $m = x_i^* + x_j^* + x_k^*$ and to the ranking that the outcomes must satisfy. If for example the ranking is $z_i \geq z_j \geq z_k$ then the problem can be written as:

$$\max_{x_i^*, x_j^*, x_k^*} w_i u(z_i) + w_j u(z_j) + w_k u(z_k) \tag{A.36}$$

$$\text{s.t. } e_j x_j - e_i x_i \leq 0 \tag{A.37}$$

$$e_k x_k - e_j x_j \leq 0 \tag{A.38}$$

$$x_i + x_j + x_k = m \tag{A.39}$$

$$x_i, x_j, x_k \geq 0^5 \tag{A.40}$$

By setting γ the Lagrangian multiplier for the budget constraint (A.39) and λ_1, λ_2 the multipliers for A.37 and A.38 respectively, we can form the Lagrangian function:

$$\mathcal{L} = w_i u(e_i x_i) + w_j u(e_j x_j) + w_k u(e_k x_k) + \gamma(m - x_i - x_j - x_k) + \lambda_1(e_i x_i - e_j x_j) + \lambda_2(e_j x_j - e_k x_k)$$

If we denote $\mathcal{L}_k$ the partial derivative with respect to k where $k = x_i, x_j, x_k, \gamma, \lambda_1, \lambda_2$, the Kuhn-Tucker conditions for this problem require:

$$\mathcal{L}_{x_i} \leq 0, x_i \geq 0, \text{and } \mathcal{L}_{x_i} x_i = 0$$

$$\mathcal{L}_{x_j} \leq 0, x_j \geq 0, \text{and } \mathcal{L}_{x_j} x_j = 0$$

$$\mathcal{L}_{x_k} \leq 0, x_k \geq 0, \text{and } \mathcal{L}_{x_k} x_k = 0$$

$$\mathcal{L}_{\gamma} \geq 0, \gamma_1 \geq 0, \text{and } \mathcal{L}_{\gamma} \gamma = 0$$

$$\mathcal{L}_{\lambda_1} \geq 0, \lambda_1 \geq 0, \text{and } \mathcal{L}_{\lambda_1} \lambda_1 = 0$$

[5] Depending on the design of the experiment this constraint can be either implemented or not. If we allow subjects to shortsell then we may let them to allocate negative amounts. If we constrain subjects from borrowing then the non-negativit constrained must hold. In this case it is possible to have boundary portfolios where all the income is allocated to one of the elements while the rest remain zero.

$$\mathcal{L}_{\lambda_2} \geq 0, \lambda_1 \geq 0, \text{and } \mathcal{L}_{\lambda_1}\lambda_1 = 0$$

which conditions can be summarised as:

$$(w_i u'(e_i x_i) - \gamma + \lambda_1 e_i) x_i = 0 \tag{A.41}$$

$$(w_j u'(e_j x_j) - \gamma - \lambda_1 e_j + \lambda_2 e_j) x_j = 0 \tag{A.42}$$

$$(w_k u'(e_k x_k) - \gamma - \lambda_2 e_k) x_k = 0 \tag{A.43}$$

$$(e_i x_i - e_j x_j)\lambda_1 = 0 \tag{A.44}$$

$$(e_j x_j - e_k x_k)\lambda_2 = 0 \tag{A.45}$$

$$x_i + x_j + x_k - m = 0 \tag{A.46}$$

Case 1: $z_i > z_j > z_k$ The problem to solve in this case is the following:

$$\max_{x_i, x_j, x_k} w_i u(z_i) + w_j u(z_j) + w_k u(z_k) \tag{A.47}$$

$$\text{s.t. } e_j x_j - e_i x_i \leq 0 \tag{A.48}$$

$$e_k x_k - e_j x_j \leq 0 \tag{A.49}$$

$$x_i + x_j + x_k = m \tag{A.50}$$

By setting γ the Lagrangian multiplier for the budget constraint and λ_1, λ_2 the multipliers for inequality constraints, we can form the Lagrangian function:

$$\mathcal{L} = w_i u(e_i x_i) + w_j u(e_j x_j) + w_k u(e_k x_k) + \gamma(m - x_i - x_j - x_k) + \lambda_1(e_i x_i - e_j x_j) + \lambda_2(e_j x_j - e_k x_k)$$

In this case, since the constrains are non binding both λ_1 and λ_2 are equal to 0, reducing the objective function to:

$$\mathcal{L} = w_i u(e_i x_i) + w_j u(e_j x_j) + w_k u(e_k x_k) + \gamma(m - x_i - x_j - x_k)$$

The optimal allocations are:

$$x_i^* = \frac{me_je_k(e_iw_i)^{1/(1-r)}}{e_je_k(e_iw_i)^{1/(1-r)} + e_ie_k(e_jw_j)^{1/(1-r)} + e_ie_j(e_kw_k)^{1/(1-r)}}$$

$$x_j^* = \frac{me_ie_k(e_jw_j)^{1/(1-r)}}{e_je_k(e_iw_i)^{1/(1-r)} + e_ie_k(e_jw_j)^{1/(1-r)} + e_ie_j(e_kw_k)^{1/(1-r)}}$$

$$x_k^* = \frac{me_ie_j(e_kw_k)^{1/(1-r)}}{e_je_k(e_iw_i)^{1/(1-r)} + e_ie_k(e_jw_j)^{1/(1-r)} + e_ie_j(e_kw_k)^{1/(1-r)}}$$

Case 2: $z_i > z_j = z_k$ When there is equality among the outcomes, this should be reflected in the weights that the decision maker puts to those. More specifically, these weights must be the same for the two equal outcomes. To show this we write down again the maximisation problem.

$$\max_{x_i,x_j,x_k} w_iu(z_i) + w_ju(z_j) + w_ku(z_k) \tag{A.51}$$

$$\text{s.t. } e_jx_j - e_ix_i \leq 0 \tag{A.52}$$

$$e_kx_k - e_jx_j = 0 \tag{A.53}$$

$$x_i + x_j + x_k = m \tag{A.54}$$

The Lagrangian now is written as:

$$\mathcal{L} = w_iu(e_ix_i) + w_ju(e_jx_j) + w_ku(e_kx_k) + \gamma(m - x_i - x_j - x_k) + \lambda_1(e_ix_i - e_jx_j) + \lambda_2(e_jx_j - e_kx_k)$$

Since $z_i > z_j$ holds with strict inequality, the respective multiplier is zero. In the case of the constraint $z_j = z_k$, the multiplier is positive. We can either solve by manipulating this multiplier, or we can plug-in the constraint to the objective function. Following the latter, and setting $z_j = z_k$ the equation reduces to:

$$\mathcal{L} = w_iu(e_ix_i) + w_ju(e_ix_i) + w_ku(e_kx_k) + \gamma(m - x_i - x_j - x_k)$$

and keeping as common factor the z_j the problem is reduced to solve for:

$$\mathcal{L} = (w_i)u(e_ix_i) + (w_j + w_k)u(e_jx_j) + \gamma(m - x_i - x_j - x_k)$$

Writing $w_j + w_k$ as $\widetilde{w}$ the objective function can be written as:

$$\mathcal{L} = (w_i)u(e_i x_i) + \widetilde{w}u(e_j x_j) + \gamma(m - x_i - x_j - x_k)$$

This $\widetilde{w}$, depending on the ranking, will be different. This is given by adding the relevant weights for the common capacities given in Table A.1s. Then it suffices to solve for x_j^*. x_k^* is equal to $\frac{e_j x_j^*}{e_k}$ and x_i will be the residual $m - x_j^* - x_k^*$. The first order conditions of the problem are:

$$L_{x_i} = 0 \Rightarrow w_i r e_i (e_i x_i)^{r-1} = \lambda \tag{A.55}$$

$$L_{x_j} = 0 \Rightarrow \widetilde{w} r e_j (e_j x_j)^{r-1} = \lambda (\frac{e_j + e_k}{e_k}) \tag{A.56}$$

Solving for the optimal allocations we obtain that:

$$x_i^* = m - x_j^* - x_k^*$$

$$x_j^* = \frac{m e_i e_k A}{e_i e_j A + e_i e_k A + eje_k (\widetilde{w} e_j e_k)^{\frac{1}{r-1}}}$$

$$x_k^* = \frac{e_j x_j^*}{e_k}$$

where $A = [(e_j + e_k) w_i e_i]^{\frac{1}{r-1}}$

Following the same methodology, it is possible to obtain the optimal allocations for the remaining two rankings in this case ($z_i = z_j > z_k$ and $z_i = z_k > z_j$).

When $z_i = z_j > z_k$, the optimal allocations are given by:

$$x_i^* = \frac{m e_j e_k A}{e_j e_k A + e_i e_k A + eie_j (\widetilde{w} e_i e_j)^{\frac{1}{r-1}}}$$

$$x_j^* = \frac{e_i x_i^*}{e_j}$$

$$x_k^* = m - x_i^* - x_j^*$$

where $A = [(e_j + e_k) w_k e_k]^{\frac{1}{r-1}}$

When $z_i = z_k > z_j$, the optimal allocations are given by:

$$x_i^* = \frac{me_je_kA}{e_je_kA + e_ie_jA + eie_k(\widetilde{w}e_ie_k)^{\frac{1}{r-1}}}$$

$$x_j^* = m - x_i^* - x_k^*$$

$$x_k^* = \frac{e_ix_i^*}{e_k}$$

where $A = [(e_i + e_k)w_je_j]^{\frac{1}{r-1}}$

Case 3: $z_i = z_j > z_k$ In this case, the solutions coincide with those of case 2. The only difference now are the weights that are attached to each of the outcomes as now the outcomes of the states that are common, are preferred to the third one.

Case 4: $z_i = z_j = z_k$ This is the trivial case where the decision maker chooses to allocate her income in such a way that no matter what is the state of the world, a fixed payoff is always guaranteed.

As the decision depends on the ranking of the outcomes, the algorithm to calculate the optimal allocation for a resolute decision maker has been programmed taking this into consideration. After forming the appropriate weights, the optimal allocation is calculated for all the possible rankings using the suitable analytical solution and the appropriate weights. These solutions then are checked if they do satisfy the ranking and the non-negative constraint and if they do, the respective utility is computed and is stored in a matrix. When all the calculation has been done, the allocation for which the utility is maximised is recovered.

The *naive* subjects fails to realize the sequential nature of the problem. Nevertheless, there is no issue of *bounded rationality*, only violation of backward induction. Consequently, a naive decision makers at the first stage behaves in the same way as a resolute does and solves the problem as if it is a static one. Then, at stage 2 receives the partial information, updates the beliefs according to the appropriate rules for the Choquet Expected Utility model and then solves the maximisation problem that involves the two remaining states. We consider the three updating rules presented earlier, the Optimistic, the Dempster-Shafer and the Generalized Bayesian updating rule. The procedure to derive the optimal allocations remains the same for the three different rules so we present only the solution in its general form that can be then easily adapted.

At stage 1 the decision maker makes the allocations based on the optimal at the present state. Then receiving the partial information there are two effects. On the one hand, the proportion of the income that was allocated to the state has not happened is lost. On the other hand, the initial beliefs on the different states are now updated based on this information. There are three events that can happen, that the ball is not i, j or k. We denote with

Consider first the case where the state of the world is "not i". The decision maker updates her beliefs for the states j and k. We denote as w'_j, w'_k the updated weights for j and k respectively. Also, the conditional allocations to j and k when the state is "not i" are denoted as $x_{j\neg i}$ and $x_{k\neg i}$. The objective now is to maximise

$$\max_{x_{j\neg i}, x_{k\neg i}} w'_j u(z_j) + w'_k u(z_k) \tag{A.57}$$

$$\text{s.t. } e_j x_{j\neg i} - e_k x_{k\neg i} \leq 0 \tag{A.58}$$

$$x_{j\neg i} + x_{k\neg i} = x^*_j + x^*_k \tag{A.59}$$

where $z_j = e_j x_{j\neg i}$ and $e_k x_{k\neg i}$ the payoffs at each case. We also write the conditional income as $m_{\neg i} = x^*_j + x^*_k$

The Lagrangian for the problem is

$$\mathcal{L} = w'_j u(e_j x_{j\neg i}) + w'_k u(e_k x_{k\neg i}) + \gamma(m_{\neg i} - x_{j\neg i} - x_{k\neg i})$$

Assuming a CRRA utility function:

$$\mathcal{L} = w'_j (e_j x_{j\neg i})^r + w'_k (e_k x_{k\neg i})^r + \gamma(m_{\neg i} - x_{j\neg i} - x_{k\neg i})$$

The first order conditions for maximisation give:

$$\mathcal{L}_{x_{j\neg i}} = 0 \Rightarrow r w'_j e_j (e_j x_{j\neg i})^{r-1} = 0$$

$$\mathcal{L}_{x_{k\neg i}} = 0 \Rightarrow r w'_k e_k (e_k x_{k\neg i})^{r-1} = 0$$

which give the optimal conditional allocations:

$$x^*_{j\neg i} = \frac{m_{\neg i} e_k (e_k w'_k)^{1/(r-1)}}{e_j (e_j w'_j)^{1/(r-1)} + e_k (e_k w'_k)^{1/(r-1)}}$$

and similarly the conditional allocation to k:

$$x^*_{k\neg i} = \frac{m_{\neg i} e_j (e_j w'_j)^{1/(r-1)}}{e_j (e_j w'_j)^{1/(r-1)} + e_k (e_k w'_k)^{1/(r-1)}}$$

or simply $x^*_{k\neg i} = m_{\neg i} - x^*_{j\neg i}$

Following the same procedure for the other two conditional states:

When the state is "not j" the allocation to i and k is given by:

$$x^*_{i\neg j} = \frac{m_{\neg j} e_k (e_k w'_k)^{1/(r-1)}}{e_i (e_i w'_i)^{1/(r-1)} + e_k (e_k w'_k)^{1/(r-1)}}$$

$$x^*_{k\neg j} = \frac{m_{\neg j} e_i (e_i w'_i)^{1/(r-1)}}{e_i (e_i w'_i)^{1/(r-1)} + e_k (e_k w'_k)^{1/(r-1)}}$$

where $m_{\neg j} = x^*_i + x^*_k$

When the state is "not k" the allocation to i and j is given by:

$$x^*_{i\neg k} = \frac{m_{\neg k} e_j (e_j w'_j)^{1/(r-1)}}{e_i (e_i w'_i)^{1/(r-1)} + e_j (e_j w'_j)^{1/(r-1)}}$$

$$x^*_{j\neg k} = \frac{m_{\neg k} e_i (e_i w'_i)^{1/(r-1)}}{e_i (e_i w'_i)^{1/(r-1)} + e_j (e_j w'_j)^{1/(r-1)}}$$

where $m_{\neg k} = x^*_i + x^*_j$

The next step is to create the weighing table. Since there are only two remaining states of the world, we need to consider only three cases, that is when $e_j x_j > e_k x_k$, $e_j x_j < e_k x_k$ and trivially $e_j x_j = e_k x_k$. In addition, we need to create this table for the three different updating rules.

Optimistic Updating Rule

The rule in its general form writes:

$$v^{OPT}_E(A) = \frac{v(A \cap E)}{v(E)}$$

Let the information revealed is "not i". In this case, the event can be written in two ways, as that the state of the world is not i on the one hand, or that the state of the world is either j or k (the union of the two available events). Applying the optimistic rule to the specific problem, the capacity for states j and k when the information revealed is "not i", is:

$$v(j|\neg i) = v(j|j\bigcup k) = \frac{v(j \cap (j\bigcup k))}{v(j\bigcup k)} = \frac{v(j)}{v(j\bigcup k)} = \frac{v_j}{v_{jk}}$$

and similarly the updated capacity for state k is:

$$v(k|\neg i) = v(k|j\bigcup k) = \frac{v(k \cap (j\bigcup k))}{v(k\bigcup k)} = \frac{v(k)}{v(j\bigcup k)} = \frac{v_k}{v_{jk}}$$

The two remaining possible states of the world are when the state is not j ($i\bigcup k$) and when the state is not k ($i\bigcup j$). In the same manner, the updated capacities are given by the following formulas:

$$v(i|\neg j) = v(i|i\bigcup k) = \frac{v(i \cap (i\bigcup k))}{v(i\bigcup k)} = \frac{v(i)}{v(i\bigcup k)} = \frac{v_i}{v_{ik}}$$

$$v(k|\neg j) = v(k|i\bigcup k) = \frac{v(k \cap (i\bigcup k))}{v(i\bigcup k)} = \frac{v(k)}{v(i\bigcup k)} = \frac{v_k}{v_{ik}}$$

$$v(i|\neg k) = v(i|i\bigcup j) = \frac{v(i \cap (i\bigcup j))}{v(i\bigcup j)} = \frac{v(i)}{v(i\bigcup j)} = \frac{v_i}{v_{ij}}$$

$$v(j|\neg k) = v(j|i\bigcup j) = \frac{v(j \cap (i\bigcup j))}{v(i\bigcup j)} = \frac{v(j)}{v(i\bigcup j)} = \frac{v_j}{v_{ij}}$$

$$v_{iE^{OPT}(A)} = \frac{v(A \cap E)}{v(E)}$$

Table A.2 Weights for the case $z_i > z_j$

	Ranking	ω_j	ω_k
1	$z_j > z_k$	$\frac{v_j}{v_{jk}}$	$\frac{v_{jk}-v_j}{v_{jk}}$
2	$z_j < z_k$	$\frac{v_{jk}-v_k}{v_{jk}}$	$\frac{v_k}{v_{jk}}$
3	$z_j = z_k$	1/2	1/2

Dempster-Shafer Updating Rule

The rule in its general form writes:

$$v_E^{DS}(A) = \frac{v((A\bigcup E^c) - v(E^c)}{(1 - v(E^c))}$$

Applying this in our problem, there are three conditional states. Let the information revealed be that the state of the world is *noti*. Then, the event E can be written as $E = \neg i$ or as the union of the two remaining states, $E = j \bigcup k$. The complementary event E^c then is simply state i, $E^c = i$. When the state of the world is not i then the decision maker needs to update her initial beliefs on the two remaining states, namely j and k. The updated capacity for j is given by:

$$v(j|\neg i) = v(j|j\bigcup k) = \frac{v(j\bigcup i) - v(i)}{1 - v(i)} = \frac{v_{ij} - v_i}{1 - v_i}$$

and for k:

$$v(k|\neg i) = v(k|j\bigcup k) = \frac{v(k\bigcup i) - v(i)}{1 - v(i)} = \frac{v_{ik} - v_i}{1 - v_i}$$

In the same way, the updated capacities for the two conditional states are given by:

When the state of the world is not j, the event is $i\bigcup k$, its complementary is $E^c = i$ and the updated capacities for states i,k are:

$$v(i|\neg j) = v(i|i\bigcup k) = \frac{v(i\bigcup j) - v(j)}{1 - v(j)} = \frac{v_{ij} - v_j}{1 - v_j}$$

$$v(k|\neg j) = v(k|i\bigcup k) = \frac{v(k\bigcup j) - v(j)}{1 - v(j)} = \frac{v_{jk} - v_j}{1 - v_j}$$

When the state of the world is not k, the event is $i\bigcup j$, its complementary is $E^c = k$ and the updated capacities for states i,j are:

$$v(i|\neg k) = v(i|i\bigcup j) = \frac{v(i\bigcup k) - v(k)}{1 - v(k)} = \frac{v_{ik} - v_k}{1 - v_k}$$

$$v(j|\neg k) = v(j|i\bigcup j) = \frac{v(j\bigcup k) - v(k)}{1 - v(k)} = \frac{v_{jk} - v_k}{1 - v_k}$$

Generalized Bayesian Updating Rule

The Generalised Bayseian updating rule in its general form is given by the following formula:

$$v_E^{GB} = \frac{v(A \cap E)}{v(A \cap E) + 1 - v(A \bigcup E^c)}$$

Applying this in our problem, there are three conditional states. Let the information revealed be that the state of the world is *noti*. Then, the event E can be written as $E = \neg i$ or as the union of the two remaining states, $E = j \bigcup k$. The complementary event E^c then is simply state i, $E^c = i$. When the state of the world is not i then the decision maker needs to update her initial beliefs on the two remaining states, namely j and k. The updated capacity for j is given by:

$$v(j|\neg i) = \frac{v(j \cap (j \bigcup k))}{v(j \cap (j \bigcup k)) + 1 - v(j \bigcup i} = \frac{v_j}{v_j + 1 - v_{ij}}$$

$$v(k|\neg i) = \frac{v(k \cap (j \bigcup k))}{v(k \cap (j \bigcup k)) + 1 - v(k \bigcup i} = \frac{v_k}{v_k + 1 - v_{ik}}$$

When the state of the world is not j, the event is $i \bigcup k$, its complementary is $E^c = i$ and the updated capacities for states i, k are:

$$v(i|\neg j) = \frac{v(i \cap (i \bigcup k))}{v(i \cap (i \bigcup k)) + 1 - v(i \bigcup j} = \frac{v_i}{v_i + 1 - v_{ij}}$$

$$v(k|\neg j) = \frac{v(k \cap (i \bigcup k))}{v(k \cap (i \bigcup k)) + 1 - v(i \bigcup j} = \frac{v_k}{v_k + 1 - v_{jk}}$$

When the state of the world is not k, the event is $i \bigcup j$, its complementary is $E^c = k$ and the updated capacities for states i, j are:

$$v(i|\neg k) = \frac{v(i \cap (i \bigcup j))}{v(i \cap (i \bigcup j)) + 1 - v(i \bigcup k} = \frac{v_i}{v_i + 1 - v_{ik}}$$

$$v(j|\neg k) = \frac{v(j \cap (i \bigcup j))}{v(j \cap (i \bigcup j)) + 1 - v(j \bigcup k} = \frac{v_j}{v_j + 1 - v_{jk}}$$

In Tables A.3A.5 the weights for the three different rules are gathered.

Table A.3 Updated Capacities Conditional to *not i*

Rule	**Ranking**	ω_j	ω_k
Optimistic	$z_j > z_k$	$\frac{v_j}{v_{jk}}$	$\frac{v_{jk}-v_j}{v_{jk}}$
	$z_j < z_k$	$\frac{v_{jk}-v_k}{v_{jk}}$	$\frac{v_k}{v_{jk}}$
Dempster-Shaffer	$z_j > z_k$	$\frac{v_{ij}-v_i}{1-v_i}$	$\frac{1-v_{ij}}{1-v_i}$
	$z_j < z_k$	$\frac{1-v_{ik}}{1-v_i}$	$\frac{v_{ik}-v_i}{1-v_i}$
Gen. Bayesian	$z_j > z_k$	$\frac{v_j}{v_j+1-v_{ij}}$	$\frac{1-v_{ij}}{v_j+1-v_{ij}}$
	$z_j < z_k$	$\frac{1-v_{ik}}{v_k+1-v_{ik}}$	$\frac{v_k}{v_k+1-v_{ik}}$

Table A.4 Updated Capacities Conditional to *not j*

Rule	**Ranking**	ω_i	ω_k
Optimistic	$z_i > z_k$	$\frac{v_i}{v_{ik}}$	$\frac{v_{ik}-v_i}{v_{ik}}$
	$z_i < z_k$	$\frac{v_{ik}-v_k}{v_{ik}}$	$\frac{v_k}{v_{ik}}$
Dempster-Shaffer	$z_i > z_k$	$\frac{v_{ij}-v_j}{1-v_j}$	$\frac{1-v_{ij}}{1-v_j}$
	$z_i < z_k$	$\frac{1-v_{jk}}{1-v_j}$	$\frac{v_{jk}-v_j}{1-v_j}$
Gen. Bayesian	$z_i > z_k$	$\frac{v_i}{v_i+1-v_{ij}}$	$\frac{1-v_{ij}}{v_i+1-v_{ij}}$
	$z_i < z_k$	$\frac{1-v_{jk}}{v_k+1-v_{jk}}$	$\frac{v_k}{v_k+1-v_{jk}}$

Table A.5 Updated Capacities Conditional to *not k*

Rule	**Ranking**	ω_i	ω_j
Optimistic	$z_i > z_j$	$\frac{v_i}{v_{ij}}$	$\frac{v_{ij}-v_i}{v_{ij}}$
	$z_i < z_j$	$\frac{v_{ij}-v_j}{v_{ij}}$	$\frac{v_j}{v_{ij}}$
Dempster-Shaffer	$z_i > z_j$	$\frac{v_{ik}-v_k}{1-v_k}$	$\frac{1-v_{ik}}{1-v_k}$
	$z_i < z_j$	$\frac{1-v_{jk}}{1-v_k}$	$\frac{v_{jk}-v_k}{1-v_k}$
Gen. Bayesian	$z_i > z_j$	$\frac{v_i}{v_i+1-v_{ik}}$	$\frac{1-v_{ik}}{v_i+1-v_{ik}}$
	$z_i < z_j$	$\frac{1-v_{jk}}{v_j+1-v_{jk}}$	$\frac{v_j}{v_j+1-v_{jk}}$

A.4.1 Sophisticated Decision Maker

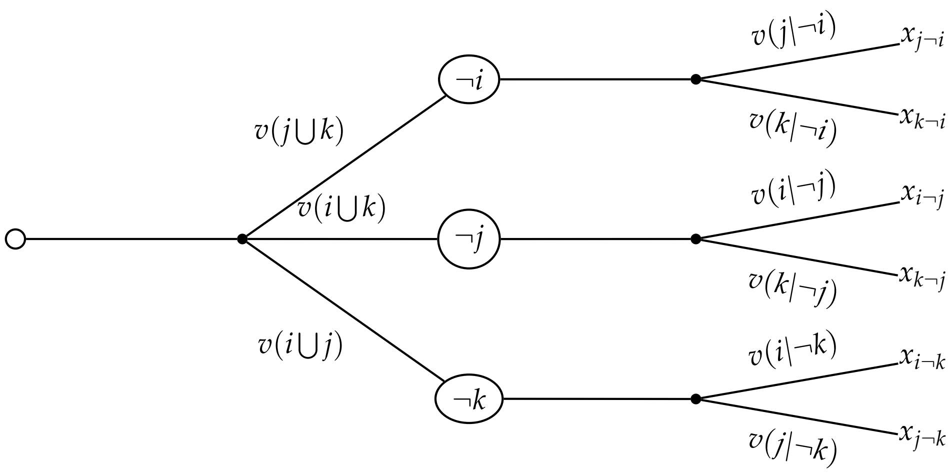

The general behaviour of the *sophisticated* type and her subsequent strategy have been presented in the main text. Here we focus on how this model can accommodate similar behaviour. Due to the fact that CEU belongs to the Rank Dependent Utility family of models, the decisions are based on weighted probabilities. In order to use the correct weights, one needs to firstly rank all the possible outcomes. In the case of a naive decision maker, this was easy to do as the problem was splited in two parts, the first involving the ranking of three outcomes (stage 1) and the second the ranking of two (stage 2-conditional). The decision tree above, shows the different stages of the decision process. At $t = 2$ all ambiguity is being resolved and the outcome is defined. This outcome is one of the conditional allocations $x_{j\neg i}, x_{k\neg i}, x_{i\neg j}, x_{k\neg j}, x_{i\neg k}, x_{j\neg k}$. It seems reasonable to expect that the allocations to the same asset conditional on different information may differ. That is to say that one may prefer the payoff of asset i when the conditional asset is *not* j to the payoff of asset i when the conditional state is *notk*. As a consequence, in order to obtain the optimal allocations one needs to consider all the possible rankings

Rewriting the problem A.60, and adjusting the compound probabilities for the respective weights ω, for the case of a Rank Dependent model, the objective function becomes:

$$\mathcal{L} = \frac{1}{2}[\omega_{j\neg i}u(e_j x_{j\neg i}(m_{\neg i})) + \omega_{k\neg i}u(e_k x_{k\neg i}(m_{\neg i})) + \omega_{i\neg j}u(e_i x_{i\neg j}(m_{\neg j})) \tag{A.60}$$
$$+\omega_{k\neg j}u(e_k x_{k\neg j}(m_{\neg j})) + \omega_{i\neg k}u(e_i x_{i\neg k}(m_{\neg k})) + \omega_{j\neg k}u(e_j x_{j\neg k}(m_{\neg k}))]$$
$$+\lambda(2m - m_{\neg i}{}^* - m_{\neg j}{}^* - m_{\neg k}{}^*)$$

For the sophisticated type, we needed to develop a computationally intensive algorithm that will be able to derive all the possible rankings among the six different outcomes and then attach the appropriate weights. If we focus for the time being at the simple case where we consider only inequalities among the possible outcomes, the possible unique permutations that we may have are 6! which means 720 possible different ways to rank the outcomes[6]. As the model is rank dependent, it is necessary for each of the rankings to attach the correct weights. Here one needs to be careful on how the weights are attached to the outcomes. For simplicity we denote all the possible rankings as a vector $k = \{x_1, x_2, x_3, x_4, x_5, x_6\}$ that corresponds to the conditional allocations, starting from *not i* to *not j*, or more analytically to the vector $\{x_{j\neg i}, x_{k\neg i}, x_{i\neg j}, x_{k\neg j}, x_{i\neg k}, x_{j\neg k}\}$. Then, a vector q contains all the compound probabilities as they were defined in Section A.2.2. Writing analytically the vector gives: $\{q_{j\neg i}, q_{k\neg i}, q_{i\neg j}, q_{k\neg j}, q_{i\neg k}, q_{j\neg k}\}$ which gives the compound

[6]For the calculation of all the permutations, the function *permutations* was used from the *R* package *GTools*.

probability that each of the outcomes will be reached.[7] For simplicity we denote them as $\{q_1, q_2, q_3, q_4, q_5, q_6\}$ As was discussed before, depending on the ranking of the outcomes, different weights are attached. Let for example the case where the following ranking is satisfied:

$$\{x_1 > x_2 > x_3 > x_4 > x_5 > x_6\}$$

and assume that the probabilities are weighted according to the way that capacities are weighted according to the CEU model or according to the following weighting function in the case of Source Choquet Expected Utility:

$$w(p) = \exp(-(-\ln(p))^{\alpha}) \tag{A.61}$$

Then the respective weights are calculated based on Table A.6 below, where w stands either for the capacity weighting or for the weighting function.

Table A.6 Weights for the Sophisticated Type at Stage 1

Outcome	**Weight**
z_1	$\omega_1 = w(q_1)$
z_2	$\omega_2 = w(q_1 + q_2) - w(q_1)$
z_3	$\omega_3 = w(q_1 + q_2 + q_3) - w(q_1 + q_2)$
z_4	$\omega_4 = w(q_1 + q_2 + q_3 + q_4) - w(q_1 + q_2 + q_3)$
z_5	$\omega_5 = w(q_1 + q_2 + q_3 + q_4 + q_5) - w(q_1 + q_2 + q_3 + q_4)$
z_6	$\omega_6 = 1 - w(q_1 + q_2 + q_3 + q_4 + q_5)$

The algorithm of the solution involves the following steps:

1. The main program provides the data for the problem (m, e_i, e_j, e_k) and the values for the parameters to estimate (p_i, p_j, r, α).
2. The compound and conditional additive probabilities are calculated.
3. A permutation is generated involving all the possible rankings among the six different outcomes.
4. The corresponding weights are calculated based on the rankings from the previous step and the weighting function.

[7] In the calculation of the probabilities, this vector is divided by 2 as this was the chance that a *not* state was announced during the experiment. This is explained extensively in Section A.2.2.

5. The conditional weights are calculated, taking into consideration the ranking at step 3 that must be satisfied.

6. For each of the permutations, the function for the optimal allocation is called with arguments the weights, the conditional probabilities and the data of the problem.

7. The allocations are checked if they satisfy the ranking. If yes, the Expected Utility is calculated and stored.

8. When all the permutations have been checked, the algorithm retrieves the allocation that generated the highest utility.

A.5 Source Choquet Expected Utility

In this model, as is the case for the CEU model, the ranking of the outcomes is important regarding the weight that is attached to each of the states of the world. Again we need to consider all the possible rankings that may exist. For a resolute decision maker, the rankings are 13 (6 strict inequalities, 6 mixed equalities and 1 case with all outcomes equal).

Table A.7 Table with all possible rankings, SCEU

	Ranking	ω_i	ω_j	ω_k
1	$z_i > z_j > z_k$	$w(p_i)$	$w(p_i \bigcup p_j) - w(p_i)$	$1 - w(p_i \bigcup p_j)$
2	$z_i > z_k > z_j$	$w(p_i)$	$1 - w(1 - p_j)$	$w(1 - p_j) - w(p_i)$
3	$z_j > z_i > z_k$	$w(p_i \bigcup p_j) - w(p_j)$	$w(p_j)$	$1 - w(p_i \bigcup p_j)$
4	$z_j > z_k > z_i$	$1 - w(1 - p_i)$	$w(p_j)$	$w(1 - p_i) - w(p_j)$
5	$z_k > z_i > z_j$	$w(1 - p_j) - w(1 - p_i - p_j)$	$1 - w(1 - p_j)$	$w(1 - p_i - p_j)$
6	$z_k > z_j > z_i$	$1 - w(1 - p_i)$	$w(1 - p_i) - w(1 - p_i - p_j)$	$w(1 - p_i - p_j)$
7	$z_i = z_j > z_k$	$w(p_i \bigcup p_j)$	$w(p_i \bigcup p_j)$	$1 - w(p_i \bigcup p_j)$
8	$z_i = z_k > z_j$	$w(1 - p_j)$	$1 - w(1 - p_j)$	$w(1 - p_j)$
9	$z_j = z_k > z_i$	$1 - w(1 - p_i)$	$w(1 - p_i)$	$w(1 - p_i)$
10	$z_i > z_j = z_k$	$w(p_i)$	$1 - w(p_i)$	$1 - w(p_i)$
11	$z_j > z_i = z_k$	$1 - w(p_j)$	$w(p_j)$	$1 - w(p_j)$
12	$z_k > z_i = z_j$	$1 - w(1 - p_i - p_j)$	$1 - w(1 - p_i - p_j)$	$w(1 - p_i - p_j)$
13	$z_i = z_j = z_k$	$1/3$	$1/3$	$1/3$

In the second stage, the beliefs are updated according to the Bayesian rule. Again, when the optimal allocation is calculated, the respective ranking should be taken into consideration. For each conditional allocation, we need to consider

Table A.8 Weights for the case *not i*

	Ranking	ω_j	ω_k
1	$z_j > z_k$	$w(\frac{p_j}{p_j+p_k})$	$1 - w(\frac{p_j}{p_j+p_k})$
2	$z_j < z_k$	$1 - w(\frac{p_k}{p_j+p_k})$	$w(\frac{p_k}{p_j+p_k})$
3	$z_j = z_k$	$1/2$	$1/2$

three rankings (two strict inequalities and 1 equality). The way that the weights are formed is shown in Table A.8.

For the sophisticated decision maker, the solution is exactly the same as in the case of CEU as this is presented in the previous section. The only difference is that the appropriate weighting function is applied.

Appendix B

Optimal Allocations for Chapter 3

In this Appendix we present the optimal allocations for both the CRRA and the CARA utility function. An algorithm is also presented for the CARA case, when the optimal solution is boundary (either zero or the total income).

B.1 Description of the Problem

The decision maker faces a transparent urn which contains balls of three different colours (we denote the colours as 1,2 and 3). This defines three states of the world $S = \{s_i, s_j, s_j\}$. Let $p_{s,i \in \{i,j,k\}}$ with $\sum_{s=i}^{S} p_s = 1$ the probability that state i happens. The probabilities are objective and known and thus the experiment can be considered as a "decision under risk" one. More specifically, the subjects are being given a series of allocation problems where they asked to allocate a given amount of experimental income (tokens) to the different colours. Each of the problem consists of an experimental income m and three exchange rates $e_{s,i \in \{i,j,k\}}$ between tokens and money, which differ at each problem. The decision maker is asked to allocate a fixed experimental income to the three different assets.

B.2 CARA Utility Function

In the experiment, the subjects are restricted to make allocations on the domain [0,m] where m the experimental income available at each problem. Consequently, we are interested in the behaviour of the CRRA utility function inside this interval. The standard Power function is defined as:

$$u(x) = \begin{cases} x^r & \text{if } r > 0 \\ \ln(x) & \text{if } r = 0 \\ -x^r & \text{if } r < 0 \end{cases}$$

The maximization problem can be written as:

$$\max_{x_i^*, x_j^*, x_k^*} p_i u(e_i x_i) + p_j u(e_j x_j) + p_k u(e_k x_k)$$

$$\text{s.t. } x_i^* + x_j^* + x_k^* = m$$

The Lagrangian writes:

$$\mathcal{L} = p_i (e_i x_i)^r + p_j (e_j x_j)^r + p_k (e_k x_k)^r + \lambda(m - x_i - x_j - x_k)$$

deriving the first order conditions:

$$\mathcal{L}_{x_i} = 0 \Rightarrow r p_i e_i (e_i x_i)^{r-1} = \lambda$$

$$\mathcal{L}_{x_j} = 0 \Rightarrow r p_j e_j (e_j x_j)^{r-1} = \lambda$$

$$\mathcal{L}_{x_k} = 0 \Rightarrow r p_k e_k (e_k x_k)^{r-1} = \lambda$$

and using the budget constraint we obtain the optimal allocations:

$$x_i^* = \frac{m(p_i e_i)^{1/(1-r)} e_j e_k}{(p_i e_i)^{1/(1-r)} e_j e_k + (p_j e_j)^{1/(1-r)} e_i e_k + (p_k e_k)^{1/(1-r)} e_i e_j}$$

$$x_j^* = \frac{m(p_j e_j)^{1/(1-r)} e_i e_k}{(p_i e_i)^{1/(1-r)} e_j e_k + (p_j e_j)^{1/(1-r)} e_i e_k + (p_k e_k)^{1/(1-r)} e_i e_j}$$

$$x_k^* = \frac{m(p_k e_k)^{1/(1-r)} e_i e_j}{(p_i e_i)^{1/(1-r)} e_j e_k + (p_j e_j)^{1/(1-r)} e_i e_k + (p_k e_k)^{1/(1-r)} e_i e_j}$$

B.2.1 CARA Utility Function

The CARA function is defined as:

$$u(x) = \begin{cases} -\exp^{-rx} & \text{if } r > 0 \\ x & \text{if } r = 0 \\ \exp^{-rx} & \text{if } r < 0 \end{cases}$$

The maximisation problem can be written as:

$$\max_{x_i^*, x_j^*, x_k^*} p_i u(e_i x_i) + p_j u(e_j x_j) + p_k u(e_k x_k)$$

$$\text{s.t. } x_i^* + x_j^* + x_k^* = m$$

The Lagrangian writes:

$$\mathcal{L} = p_i e^{-re_i x_i} + p_j e^{-re_j x_j} + p_k e^{-re_k x_k} + \lambda(m - x_i - x_j - x_k)$$

deriving the first order conditions:

$$\mathcal{L}_{x_i} = 0 \Rightarrow r p_i e_i e^{-re_i x_i} = -\lambda$$

$$\mathcal{L}_{x_j} = 0 \Rightarrow r p_j e_j e^{-re_j x_j} = -\lambda$$

$$\mathcal{L}_{x_k} = 0 \Rightarrow r p_k e_k e^{-re_k x_k} = -\lambda$$

$$r p_i e_i e^{-re_i x_i} = r p_j e_j e^{-re_j x_j} \tag{B.1}$$

$$r p_j e_j e^{-re_j x_j} = r p_k e_k e^{-re_k x_k} \tag{B.2}$$

$$r p_i e_i e^{-re_i x_i} = r p_k e_k e^{-re_k x_k} \tag{B.3}$$

Finally substituting to the budget constraint we obtain the optimum value for x_i:

$$x_i + \frac{\ln(\frac{p_j e_j}{p_i e_i})}{re_j} + \frac{e_i x_i}{e_j} + \frac{\ln(\frac{p_k e_k}{p_i e_i})}{re_k} + \frac{e_i x_i}{e_k} = m$$

$$x_i^* = \frac{me_je_k - \left[\ln(\frac{p_ke_k}{p_ie_i})e_j + \ln(\frac{p_je_j}{p_ie_i})e_k\right]/r}{e_ie_k + e_je_k + e_ie_j}$$

and similarly we obtain:

$$x_j^* = \frac{me_ie_k - \left[\ln(\frac{p_ke_k}{p_je_j})e_i + \ln(\frac{p_ie_i}{p_je_j})e_k\right]/r}{e_ie_k + e_je_k + e_ie_j}$$

$$x_k^* = \frac{me_ie_j - \left[\ln(\frac{p_ie_i}{p_ke_k})e_j + \ln(\frac{p_je_j}{p_ke_k})e_i\right]/r}{e_ie_k + e_je_k + e_ie_j}$$

As the CARA utility function is not bounded, it may be the case where the subjects would like to short-sell and thus to allocate negative amounts to specific colours. As the experimental design did not allow for this kind of behaviour and forced the participants to allocate anything between zero and the total income, we need to take into consideration these extreme cases. For example, if a decision maker wants to allocate a negative amount to one of the available options, based on her preferences, maybe it is optimal to allocate zero in this colour and concentrate in the remaining options. This kind of optimisation thinking must be taken into consideration when the optimal allocations are calculated. To tackle this issue we extend the current algorithm with some additional checks. Using the formulas derived above, we are able to control if the optimal decision requires negative allocations. In this case, we can distinguish 2 different extremes. On the one hand there is the case where the total income is allocated only to one of the available options while on the other, the case where nothing is allocated to one of the options and the total income is devoted to the remaining two available colours. The next table summarizes the possible case we may observe and the respective formulas that provide the optimal allocations.

The algorithm first finds the optimal allocation for the unconstrained problem. If the solution satisfies the non-negativity constraints, then the solution is acceptable and there is no need for further investigation. If there are negative, the algorithm calculates the optimal allocation based on the cases described on Table B.1. The values are again checked if they satisfy the constraints. If the answer is positive, the respective expected utility is calculated and stored along with the respective allocation. When this process has been applied to all the sub-cases, the

Table B.1 Extreme Allocations

			x_i^*	x_j^*	x_k^*
$x_i^* < 0$	$x_j^* \geq 0$	$x_k^* \geq 0$	0	$\frac{mx_k - \frac{\ln(p_k * x_k)}{(p_j x_j)}/r}{x_j + x_k}$	$\frac{m * x_j - \frac{\ln(p_j x_j)}{(p_k x_k)}/r}{x_j + x_k}$
$x_i^* > 0$	$x_j^* < 0$	$x_k^* < 0$	m	0	0
$x_i^* \geq 0$	$x_j^* < 0$	$x_k^* \geq 0$	$\frac{mx_k - \frac{\ln(p_k * x_k)}{(p_i x_i)}/r}{x_i + x_k}$	0	$\frac{m * x_i - \frac{\ln(p_i x_i)}{(p_k x_k)}/r}{x_i + x_k}$
$x_i^* < 0$	$x_j^* \geq 0$	$x_k^* < 0$	0	m	0
$x_i^* \geq 0$	$x_j^* \geq 0$	$x_k^* < 0$	$\frac{mx_j - \frac{\ln(p_j * x_j)}{(p_i x_i)}/r}{x_i + x_j}$	$\frac{m * x_i - \frac{\ln(p_i x_i)}{(p_j x_j)}/r}{x_i + x_j}$	0
$x_i^* < 0$	$x_j^* < 0$	$x_k^* \geq 0$	0	0	m

algorithm retrieves the allocation that generates the highest possible utility.

Appendix C

The Experimental Instructions - Experiment in Chapter 2

Instructions

Preamble

Welcome to this experiment. These instructions are to help you to understand what the experiment is about and what you are being asked to do during it. The experiment gives you the chance to earn a considerable amount of money, which will be paid to you in cash after you have completed the experiment. The payments described below are in addition to a participation fee of £2.5 that you will be paid independently of your answers.

The Bingo Blower

At the back of this laboratory you will see a Bingo Blower which also projected in the two screens. In this, as you will see, there are balls of three different colours - Pink, Blue and Yellow - which are being blown around inside the Blower. Please familiarize yourself with this device as your payments depend on the composition of the Bingo Blower. You can observe the screens during the experiment and you are welcomed at any time to come inside the room where the Bingo Blower is and have a closer look. After you have responded to the various questions in the experiment, one of the 60 questions will be chosen at random. Then, you will go to the Bingo Blower and activate a mechanism which will expel one ball (depending on the specification-see below) from the Bingo Blower. The colour of this one ball, combined with your answer to the question picked at random, will determine your payment for taking part in this experiment - as we describe below.

The Questions in the Experiment

You will be asked a total of 60 questions. Each question has two stages. A Three-Colour Fist-Stage and one Two-Colour Second-Stage. You are asked to allocate a given quantity of tokens among the three colours in the Bingo Blower. Then you receive some piece of information and you are asked to make an new allocation between two colours.

The First Stage

In the first stage, you will be given a quantity of tokens and you will be asked to allocate them among the three colours in the Bingo Blower: that is, Yellow, Pink and Blue. You will also be told the exchange rate between tokens and money for each colour.

The Second Stage

In the second stage you will be told that a ball is drawn but you will not learn its colour. You will learn that it is not of colour x (where x is Yellow, Blue or Pink). Then you will be asked to allocate tokens between the two remaining colours in the Bingo. The amount of available tokens that you have in the second stage depends on how much you allocated to each colour at the first stage. The ball that is drawn is the 'winning colour' and your final payment depends on the amount you allocated to the 'wining colour'. For example, you make your allocation at stage 1. You learn that a ball is drawn and this ball is NOT Pink. Then you are asked to make an allocation to the two remaining colours, thus Yellow and Blue. When you receive the information that the ball is NOT Pink, you lose all the tokens that you have allocated to Pink at stage 1. So the number of tokens you have in the second stage is equal to the sum of tokens you allocated in the first round to Yellow and Blue.

Payment

At the end of the experiment one of all the 60 questions will be played for real. You will be asked to randomly choose one number from 1 to 60 which will be the problem according to which you will be paid. The computer will recover your allocation at this problem and what was the information you received (e.g. not Pink). You will then activate the Bingo Blower and a ball will be drawn. The colour of this ball will determine the 'winning colour' (if the ball is the same colour to the one that you were informed that is not, you continue drawing balls till a different colour will be drawn). Your payment is the amount you allocated at that specific problem to the 'winning' colour expressed in money (multiplied by respective exchange rate of this colour).

Example of a Question

Suppose the question asks you to allocate 60 tokens among Pink, Blue and Yellow. Suppose the exchange rate for Pink is 0.5, the exchange rate for Blue is 1 and the exchange rate for Yellow is 0.4. Observing the Bingo Blower and the likelihood of each colour to be drawn, you decide to allocate: Pink: 20 tokens Blue: 25 tokens Yellow: 15 tokens

Then you learn that a ball is drawn and this ball is NOT Pink. Now you have to allocate the remaining tokens between Blue and Yellow. Since the ball is NOT Pink, what you allocated at the first stage to Pink is not anymore available. This means that now you have 40 tokens (25+15) to allocate between Yellow and Blue balls. In this second stage, if you allocate 20 tokens at Blue and 20 tokens at Yellow the implying payment is: £20 (1*20) if the 'winning' colour is Blue £8 (0.4*20) if the 'winning' colour is Yellow

PowerPoint Presentation

To ensure that you understand these instructions, in your computer there is a PowerPoint presentation. Then you should call over an experimenter, who will check that you are happy before initiating the experiment. After the experiment is over, we will ask you to sign a receipt and then you will be free to go. Thank you for your participation.

How Long the Experiment Will Last

Note that the experiment will take at most 2 hours as you are given 60 seconds for each stage. It is clearly in your interests to be as careful as you can when you are answering the questions.

Konstantinos Georgalos
John Hey
June 2013

Appendix D

The Experimental Instructions - Experiment in Chapter 3

Instructions

Preamble

Welcome to this experiment. These instructions are to help you to understand what you are being asked to do during the experiment and how you will be paid. The experiment is simple and gives you the chance to earn a considerable amount of money, which will be paid to you in cash after you have completed the experiment. The payment described below is in addition to a participation fee of £2.50 that you will be paid independently of your answers.

The Experiment

In this experiment, you will be presented with a total of 72 problems. Each problem is of the same form. You will be asked to allocate a total of 100 tokens to three different colours: green, blue and red. The problems will vary in two ways: first the money value of the tokens allocated to each of the three colours; second, the chances of green, blue and red occurring. After you have responded on all 72 problems, one of them will be chosen at random and played out. 'Playing out' a problem means that one of green, blue or red will occur (in a way we describe below) with chances as specified in that problem, and you will be paid the money value of the number of tokens that you allocated to that of green, blue and red that occurred.

Specifying the Chances

Each problem will specify the chances of green, blue and red occurring. These may differ from problem to problem but they will always be something out of 10. For example a problem may specify the chance of green occurring as 5 out of 10, of blue occurring as 3 out of 10, and of red occurring as 2 out of 10. If this problem were to be played out at the end of the experiment then you would put 5 green balls, 3 blue balls and 2 red balls into an opaque bag. The bag would then be shaken and you would be asked to draw a ball at random (and without looking into the bag) out of the bag. Clearly the chance of drawing a green ball is 5 out of 10, of drawing a blue ball 3 out of 10, and of drawing a red ball 2 out of 10.

Your Payment

As already noted, at the end of the experiment one of the 72 problems will be randomly selected by your randomly choosing one lottery ticket from a set of tickets numbered from 1 to 72; the number on the ticket randomly chosen will determine the problem to be played out. At this point, the experimenter will recall your allocation on that particular problem, the money values of tokens allocated to green, blue and red, and also the chances of green, blue and red in that problem. You will then compose the bag as described in the paragraph above, and then you will draw one ball at random (as described above) out of the bag. The money value of your allocation to the colour drawn will be your payment. Note that you cannot determine the colour of the ball drawn, but you do know the chances of it being green, blue or red.

Example of a Problem

Imagine a problem where the chances of green, blue and red are specified as 5 out 10, 3 out of 10 and 2 out of 10 respectively. You have 100 tokens to allocate to green, blue and red. Suppose that the money values in this problem are: green: 1 token has value 30p blue: 1 token has value 40p red: 1 token has value 40p

Suppose that you make the following allocation: green: 60 tokens blue: 10 tokens red: 30 tokens

Then you will put 5 green, 3 blue and 2 red balls into a bag which will then be shaken and you randomly draw a ball. Then your payment would be: 60 times

30p = £18.00 if the ball you draw is green 10 times 40p = £4.00 if the ball you draw is blue 30 times 40p = £12.00 if the ball you draw is red

The Experimental Interface

The computer will tell you in each problem the money implications of any allocation. You will see a screen like the one below. This represents the problem (and a possible response) on the problem discussed above.

At the top left of the screen you are told the problem number, the numbers of balls of the three different colours in the bag if this problem were to be played out, and the exchange rate between tokens and money for the three colours. The triangle is the crucial thing. When you move the cursor around within the triangle your allocation of the 100 tokens changes. The coloured areas within the triangle represent the proportional allocations, and the coloured numbers specify the allocations to the respective colours. The picture at the bottom right indicates the amounts implied by your allocations, written in the appropriate colours underneath the figure and represented by the heights of the coloured bars; the widths of the coloured bars represent the chance of the colours being realised in that problem. You should move your cursor around within the triangle until you reach your preferred allocation. When you have found your desired allocation, you should double click on your mouse and then confirm your decision (or cancel it and select again if you change your mind). When you click on 'Confirm' you will move onto the next problem. Remember that there are 72 of them.

Timing

You will not be able to confirm your decision on any problem until at least 15 seconds have elapsed. You will be given a maximum of 60 seconds to decide; if you have not confirmed your decision by the time 60 seconds have elapsed, it will be assumed that you wish to allocate your tokens (as near as possible) equally between the three colours.

PowerPoint Presentation

To ensure that you understand these instructions, we will display a PowerPoint presentation of the Instructions. In this presentation, the interface where

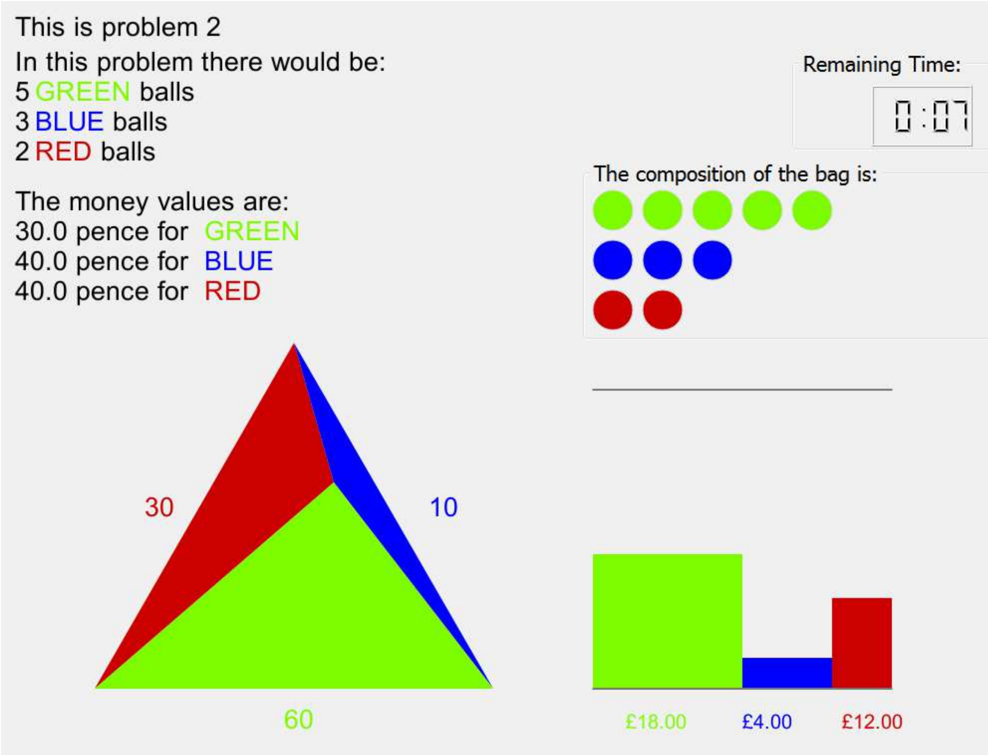

Figure D.1 Screenshot of the Experimental Framework

you will be able to read the questions and submit your answers is explained. During the session, experimenters will be around so please feel free to ask them for clarifications. When you finished watching the presentation please call over an experimenter, who will check that you are happy before initiating the experiment and will activate your computer. After the experiment is over, we will ask you to sign a receipt. After you will get paid and then you will be free to go.

How Long the Experiment Will Last

You will be given a maximum of 60 seconds to determine your allocation in each problem, so the experiment will last at most one-and-a-half hours. It is clearly in your interests to be as careful as you can when you are answering the questions.

Thank you for your participation

Xueqi Dong
Konstantinos Georgalos
John Hey
June 2014