Dynamic Decision Making under Ambiguity: An Experimental Investigation

Konstantinos Georgalos*

December 13, 2020

Abstract

Neoclassical economic theory assumes that whenever agents tackle dynamic decisions under ambiguity, preferences are represented by the Subjective Expected Utility (SEU) model and prior beliefs are updated according to Bayes rule, upon the arrival of partial information. Nevertheless, when one considers non-neutral ambiguity attitudes, either the axiom of *dynamic consistency* or of *consequentialism* should be relaxed. Using data from an economic experiment on dynamic choice under ambiguity, we study which of the two rationality axioms people violate, along with the question of whether this violation is part of a conscious planning strategy or not. The combination of the two, allows us to classify non-SEU subjects to three behavioural types: resolute, naïve and sophisticated. The hypothesis of Bayesian updating is rejected for more than half of the experimental population. For ambiguity non-neutral subjects, we find that the majority are sophisticated, a few are naïve and very few are resolute.

JEL classification: C91, D81, D83, D90

Keywords: Ambiguity, Subjective Beliefs, Dynamic Consistency, Consequentialism, Updating, Experiment

---

*Lancaster University Management School, Department of Economics, LA1 4YX, Lancaster, U.K.
✉ k.georgalos@lancaster.ac.uk, ☎ +44 (0)15245 93170
I am grateful to John D. Hey, Glenn Harrison and Ivan Paya for their valuable comments and guidance. I am also thankful to Xueqi Dong and Hushuan Li for their assistance during the experimental sessions. This research was funded by a Research and Impact Support Fund awarded by the Department of Economics at the University of York (RIS 39). The financial aid of the Greek Scholarships Foundation (IKY) is gratefully recognised.

1 Introduction

Underlying much of economic theory are three key assumptions. These are that economic agents: (1) use probabilities to describe risky and ambiguous situations; (2) behave in a dynamically consistent way; and (3) update probabilities according to Bayes rule, upon the arrival of partial information. Subjective Expected Utility theory (SEU, Savage, 1954) binds these three assumptions together in a logical and intellectually satisfying manner. Nevertheless, since the seminal thought experiments proposed by Ellsberg (1961), challenging the first assumption, a vast literature of theoretical models emerged, aiming to accommodate Ellsberg-type preferences[1]. The direct consequence of this was the rapid development of a large body of experimental work, that either tests the attitudes towards ambiguity, or performs *horse-race* comparisons to identify the model that best describes data[2].

However, as it is highlighted in Gilboa and Schmeidler (1993), if one wants to confirm the theoretical validity of any model of decision making under uncertainty, this model should be able to successfully cope with the dynamic aspect of the choices. When SEU is extended to its dynamic dimension, the independence axiom (often called the "sure thing principle") is equivalent to two rationality axioms, namely *dynamic consistency* (DC) and *consequentialism* (C), along with other conventional assumptions. DC requires that the *ex-ante* preferences coincide with the *ex-post* ones, while C dictates that past decisions play no role and only available options matter[3]. Ghirardato (2002), provides the elegant result that when both DC and C are satisfied, preferences are represented by SEU and the agent's beliefs are updated according to Bayes rule[4]. However, given that most of the non-SEU models relax the independence

[1]See among others Gilboa and Schmeidler (1989), Schmeidler (1989), Tversky and Kahneman (1992), Ghirardato et al. (2004), Klibanoff et al. (2005), Maccheroni et al. (2006), Gajdos et al. (2008), Siniscalchi (2009). For an extensive review of the models see Etner et al. (2012).

[2]See Halevy (2007), Hayashi and Wada (2010), Hey et al. (2010), Abdellaoui et al. (2011), Charness et al. (2013), Hey and Pace (2014), Ahn et al. (2014), Stahl (2014), Baillon and Bleichrodt (2015). For a review of the main results see Trautmann and van de Kuilen (2015).

[3]Consequentialism was first proposed in Hammond (1988) and it requires that the conditional preferences to remain unaffected by the outcomes outside the conditional events. Representing the dynamic problem with a decision tree, consequentialism is satisfied when the decision maker 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.

[4]Klibanoff and Hanany (2007) claim that dynamic consistency is the primary justification for Bayesian updating and under the view that Bayesian updating should be taken as given, DC comes "for free" under Expected Utility.

axiom, modelling dynamic choice requires the theoretician to abandon either DC or C and consequently, to abandon Bayes rule.

Al-Najjar and Weistein (2009), classify theories to four different categories, depending the assumptions they make on how decision makers (DMs) tackle dynamic problems and update their beliefs upon the reception of partial information. The first category includes theories that abandon DC and are labeled as "naïve updating" theories, since it is not necessary for the decisions at the present to take into consideration future preferences. This includes Gilboa and Schmeidler (1993), Pires (2002), Wang (2003), Eichberger et al. (2007) and Eichberger et al. (2010)[5]. In this approach, each of the stages is faced independently of the other, strategy that may lead to dynamic inconsistencies and dominated results. The second category, includes theories that require the DM to behave in a sophisticated way, thus violating DC. This approach is mainly represented by Siniscalchi (2011), who does not assume any particular preference functional or update rule. The idea is based on the notion of *consistent planning*, where the *ex-post* preferences are taken into account when the *ex-ante* choices are made. An alternative way to model dynamic choice involves the relaxation of C. This family of models proposes the use of a set of distorting updating rules that ensure DC. This includes Klibanoff and Hanany (2007), Hanany and Klibanoff (2009) and Klibanoff et al. (2009), who have axiomatised and extended few of the most commonly used ambiguity models to their dynamic version. Finally, there is a category of models in the literature that maintains both C and DC in the framework of multiple-priors representation of beliefs. To this end, these models require the restriction of information sets and allow the updating only of the set of beliefs that does not reverse the *ex-ante* choices based on the *rectangularity* condition. A representative model of this approach is presented in Epstein and Schneider (2003)[6].

This classification, resembles Machina (1989), who defines four different types of DMs in dynamic choice under risk: the so called α*-people*, are dynamically consistent agents who maximise EU preferences, the β*-people* who are non-EU agents and apply consequentialism, acting in a dynamically inconsistent way (myopic behaviour),

[5] Ozdenoren and Peck (2008) in a game theoretical framework, show that violating DC is the rational course of action, when suspicion is perceived regarding the composition of the Ellsberg urn.

[6] An exhaustive review of the theoretical literature is beyond the scope of this study. Al-Najjar and Weistein (2009), Klibanoff et al. (2009) and Siniscalchi (2011), provide excellent reviews of the various approaches on modelling dynamic preferences under ambiguity.

the γ-*people* who are non-EU agents but are dynamically consistent and finally, the δ-*people* who are also non-EU, are characterised as *sophisticated* and satisfy consistent planning[7].

In this study, we use data from an experiment on dynamic decision making and we classify subjects to four behavioural types, the SEU maximiser, the *naïve*, the *resolute* and the *sophisticated* type which correspond to the β-*people*, γ-*people* and δ-*people* respectively.

Our aim is to understand how people behave in a dynamic decision problem under ambiguity where decisions are made before and after the resolution of some uncertainty. Using two-stage allocation questions, with partial information revealed in an interim stage, this study aims to investigate three main questions: first, do people behave according to the predictions of the SEU model and therefore, update beliefs in a Bayesian way? Second, when people deviate from SEU, which of the two rationality axioms do they violate? Third, when subjects violate the axioms of SEU, are they aware of this violation? In other words, is this violation the consequence of a conscious planning strategy? The latter allows the classification of subjects to the various behavioural types.

Surprisingly, experimental studies of dynamic decision making under ambiguity are quite scarce. Cohen et al. (2000), in a non-incentivised experiment, study the descriptive validity of the main two updating rules that have been axiomatised for the multiple-priors family, the *Maximum Likelihood Updating* (MLU) rule and the *Full Bayesian Updating* (FBU) rule. Using a design based on the dynamic Ellsberg urn, they confirm the Ellsberg type behaviour and show that the FBU rule is applied more often. They assume *separability* (an assumption close to C), which does not allow for a direct test of which of the two axioms a subject satisfies. Dominiak et al. (2012) use a similar design as in Cohen et al. (2000). They test whether subjects satisfy DC or C, providing evidence of extensive violation of DC, whilst finding supportive evidence for the FBU rule. Finally, Georgalos (2019), using a set of two-stage allocation questions with partial revelation of information, and assuming naïve updating, tests the predictive capacity of various dynamic ambiguity models, finding support for the Tversky and

[7] A similar classification has been also applied in a hyperbolic discounting context (O'Donoghue and Rabin, 1999) and later in Hey and Panaccione (2011) and Barberis (2012), in the context of dynamic decision making under risk.

Kahneman (1992) Prospect Theory specification.

Georgalos (2019) assumes consequentialism and estimates various models and updating rules. Although the present study will use the same data set as in Georgalos (2019), we extend their analysis in two ways: first, we test which of the two axioms is violated by ambiguity non-neutral agents. Second, by defining various behavioural rules in planning strategies, we investigate whether subjects are aware of this violation and whether they take it into consideration when they make choices.

Overall, we find substantial heterogeneity in behaviour. Only half of our experimental population behaves according to SEU. For the ambiguity non-neutral subjects, the majority are best described by the sophisticated type, few by the naïve and a small minority by the resolute. We also find extensive violations of dynamic consistency.

The rest of the paper is organised as follows: section 2 presents the decision task, section 3 presents the theoretical framework and the different planning strategies, section 4 summarises the assumptions for the statistical model and the estimation method, while section 5 presents the results. We then conclude.

2 Data and the Portfolio Choice Problem

In this paper we use experimental data to estimate various preference functionals. We elicit beliefs and strategies based on a *revealed preference* argument. To achieve this, we use data from a series of two-stage portfolio allocation questions, in an experimental design inspired by Loomes (1991). This allocation procedure has the potential to provide more informative data and has been generally applied in the literature in various contexts[8].

The two-stage portfolio allocation task consists of three, payoff relevant, mutually exclusive states of nature. Each state s, corresponds to state-contingent *Arrow* security, the return of which equals e_s if state s occurs and 0 otherwise. e_s is the rate of return of asset s (henceforth exchange rate) for every unit of income allocated to this asset.

[8] Studies that use allocation problems include Choi et al. (2007) in a portfolio choice experiment under risk, Charness and Gneezy (2010) studying portfolio choices, Hey and Panaccione (2011) on dynamic decision making under risk, Ahn et al. (2014) in a portfolio choice experiment under ambiguity, Hey and Pace (2014) comparing different static models of choice under ambiguity and Loomes and Pogrebna (2014) studying individual risk attitudes. See Loomes and Pogrebna (2014) for an extensive discussion on the allocation procedure.

Ambiguity is represented with the aid of a transparent and non-manipulable device, a Bingo Blower. Inside the transparent Bingo Blower, there were 20 balls in total, of three different colours (red, blue and yellow), representing the three potential states of the world. The balls were in continuous motion and the actual composition of the Bingo Blower consisted of 4 blue (20%), 6 red (30%) and 10 yellow (50%) balls out of the total 20. The advantage of this device is that although there exist objective probabilities, known only to the experimenters, the subjects lack this information and they somehow form subjective beliefs, regarding the probabilities of the three states of the world, that may not be known to them in a conscious way. That is, while they can observe that there is at least one ball of each colour (lower bound probability), it is almost impossible to be able to identify the exact composition of the Bingo Blower, generating, in this way, genuine ambiguity. During the experiment, the subjects could observe the physical Bingo Blower placed in the middle of the lab and they could also consult live streaming of the blower projected in two large screens in the front of the lab.

Each allocation question consists of an experimental income m, expressed in tokens, and an exchange rate e_s for each of the possible states of the world, with s being either blue, red or yellow. At the first stage, the subjects were asked to allocate their experimental income between the three colours, knowing that at the end of the second stage, one ball would be extracted from the Bingo Blower, and the payoff would be the product between the tokens allocated to this colour and its respective exchange rate. Before that, there was an interim stage, where the subjects would obtain partial information regarding the ball that was extracted, but they would not learn the actual colour of the ball. For example, the subjects would learn that "the ball is not blue". Then all the tokens that were allocated to the blue colour were lost and the subject was given the opportunity (if she wished so) to re-allocate the remaining endowment to the two available colours, red and yellow. The subjects knew that during the interim stage, they would receive the information that the colour is not blue, yellow or red with equal chances.

At $t = 2$, all ambiguity would be resolved, the actual state of the world would be revealed and the DM would be paid the state-contingent dividend. A crucial aspect of the experimental design is that the revealed preference methodology was based on the *strategy* method, that is, but the only time the period t=2 was carried out was after the experiment when the question that would be played for real was presented. Realising $t = 2$ for every question, and therefore allowing continuous sampling from the Bingo Blower, could potentially generate *learning effects* regarding the actual probability distribution, thus transforming the problem into a risky situation (ambiguity/uncertainty is eliminated) with objective probabilities[9]. As the objective of the experiment was to estimate belief updating models, rather than belief learning ones, participants were informed that no actual draws will take place during the experiment, but instead, hypothetical draws, and subsequently hypothetical partial information would be announced to subjects.[10].

The same task was repeated for a total of sixty, independent allocation questions, where each question involved a varying endowment from 9 to 110 tokens, while the exchange rates for the three colours were varied between 0.1 to 1.8[11]. The subjects had to allocate all their endowment in each question.

Subjects were paid according to a random incentive mechanism, where one of the allocation questions was played for real. There are data from 58 subjects (see Georgalos 2019 for analytical details on the design).

3 Theoretical Framework and the Different Types

In this section we present the latent structural models of decision making that we fit to our data, as well as the various behavioural types of DMs. First we provide the formal definitions of the axioms needed to characterise the different types, namely dynamic consistency, consequentialism and the principle of consistent planning. Using a theo-

[9]See for example Trautmann and Zeckhauser (2013), Ert and Trautmann (2014) and Baillon et al. (2018).

[10]This form of hypothetical signals has been previously applied in the literature in Griffin and Tversky (1992) and Kraemer and Weber (2004). The computer was programmed to draw i.i.d. virtual balls from a uniform distribution. Then it would announce that the ball is not one of the two remaining colours with equal chances. For instance, if a ball was blue, it would announce that the ball is not red with probability 50% and not yellow with the residual probability. This procedure was communicated to the subjects. The information that the subjects received on average was “not blue” 34.54%, “not yellow” 34.36% and “not red” 31.09%.

[11]The questions have been chosen after extensive Monte Carlo simulations that would ensure three issues: (1) that for a simulated dataset using a given set of parameter values, it is possible to estimate (recover) the value of the actual parameters; (2) that it is possible to identify between the different specifications and; (3) that our estimation programs work efficiently. See section 4 for details on the econometric analysis.

retical framework similar to Savage (1954), we define a state space S, which includes all the possible states of the world (in our framework $S = \{R, B, Y\}$ for the red, blue and yellow assets), and a set of outcomes C, which includes the payoffs from the portfolio. An act f is a mapping from the state space S, to the set of outcomes C, which in our framework corresponds to an allocation. An event E is a subset of S. In the general case, an act $f_E g$ assigns the outcome $f(s)$ to each state of nature $s \in E$, and the outcome $g(s)$ to each state $s \in S \backslash E$. A decision maker is endowed by preferences $\succsim$ over the set of all possible acts $\mathcal{F}$. After learning that an event E has occured, the decision maker updates her beliefs and constructs conditional preferences represented by $\succsim_E$. Also, let $u : \mathbb{R} \rightarrow \mathbb{R}$ be a standard utility function that satisfies the usual assumptions of being twice differentiable, strictly increasing and strictly concave.

Dynamic Consistency. *An agent satisfies dynamic consistency (DC) if for any non-empty event E and all acts $f, g \in \mathcal{F}$, such that $f(s) = g(s)$ for each $s \in S \backslash E$, $f \succsim g$ implies $f \succsim_E g$.*

While in a pairwise choice context, DC dictates the lack of preference reversals, in our allocation context, DC should guarantee that the *ex-ante* allocations between two assets should be identical to the *ex-post* allocations, conditional on the partial information that the event E provides.

Consequentialism. *An agent satisfies consequentialism (C) when for every non-null event E and all acts $f, g \in \mathcal{F}$, $f(s) = g(s)$ for each $s \in E$ implies $f \sim_E g$.*

More specifically, this axiom is divided in two parts: that no weight is placed on the consequences of acts that are not available any more, and that the conditional preferences depend only on the information provided by the conditioning event E. Al-Najjar and Weistein (2009) refer to this type of updated preferences as *fact-based* updated. In the context of the dynamic portfolio choice problem, consequentialism implies that previous allocations to the state that is no longer available, play no role on the conditional preference between assets. In addition, it implies that inconsequential characteristics of the dynamic problem, such as the *ex-ante* optimal plan or various feasibility constraints at previous stages of the problem, have no influence to the conditional preferences. Machina (1989) describes consequentialism as if the decision maker is snipping off or ignoring any part of the tree that can no longer be reached. In that sense, the DM is forward looking and history has no effect to subsequent choices.

Finally, the last definition is required for those DMs who proceed to violate DC, but are aware of this violation. It is based on the notion of *consistent planning*, first introduced in Strotz (1955-56) and extended in Siniscalchi (2011).

Consistent Planning. *An agent adopts the consistent planning (CP) strategy, if at each decision node, the best plan among those that will be actually followed is chosen.*

This concept borrows elements from the game theoretical literature, where the dynamic problem is represented by a game played by multiple selves of the same individual. The DM applies *backward induction* and her planning strategy requires to first consider the terminal choice node of a decision tree and choose the optimal course of action at this point. Then, by "folding back", she calculates the optimal choice in the previous nodes, taking into consideration her future preferences. Siniscalchi (2011) formally axiomatises this concept for dynamic choice under ambiguity by deriving *ex-ante* conditional preferences over decision *trees* rather than over *acts*. We next describe the behaviour for each of the specifications that we consider. We start by presenting the benchmark model of SEU with Bayesian updating and then, we subsequently relax the assumption of ambiguity neutral attitudes. We present the strategies assuming a generic form regarding the utility representation.

3.1 Subjective Expected Utility

The DM is assumed to hold a unique set of subjective, additive priors $\pi = \{\pi(R), \pi(B), \pi(Y)\}$ regarding the three possible states of the world such that $\pi(R) + \pi(B) + \pi(Y) = 1$. As already highlighted, a convenient feature of the SEU model is that the DM satisfies both DC and C. Consequently, the beliefs of the agent are updated according to the Bayes rule which ensures that the *ex-ante* allocation coincides with the *ex-post*. Hence, it suffices to solve the problem as if it was a static one with three possible states of the world. The objective of the DM is to calculate the optimal portfolio $X = (x_R, x_B, x_Y)$, based on her subjective beliefs, that maximises the expected utility according to the utility function $u(.)$, subject to the budget and the non-negativity constraints. The optimal allocation is given by maximising:

$$\max_X \pi(R)u(e_R x_R) + \pi(B)u(e_B x_B) + \pi(Y)u(e_Y x_Y)$$
$$\text{s.t. } x_R + x_B + x_Y = m$$

where x_s is the amount of tokens allocated to asset s and e_s is its exchange rate ($e_s \times x_s$ determines the payoff of asset s). By definition, a DM that holds additive subjective beliefs is characterised by a *neutral* attitude towards ambiguity.

3.2 The α-Maxmin Model

Here we relax the assumption of additive beliefs and we introduce non-neutral ambiguity attitudes assuming that the DM has α-Maxmin preferences (α-MEU, Ghirardato et al., 2004). In this model the agent believes that the true probabilities over the state space lie within a continuous, closed and convex set of subjective priors Π (multiple-priors representation). This set includes all the possible scenarios regarding the future states of the world, in the form of subjective probability distributions (beliefs). Figure 1 illustrates this set Π using a two-dimensional unit simplex (known as the Marschak-Machina Triangle[12]) where the probability that the state of the world is R (Y) is represented in the horizontal (vertical) axis. Assuming that there exist non-zero low bounds of the DM's subjective beliefs ($\underline{\pi}(R), \underline{\pi}(B), \underline{\pi}(Y)$), we are able to draw the interior triangle, the size of which illustrates the degree of *ambiguity perception* of the agent. When this interior triangle shrinks to a single point, then all ambiguity vanishes and the model reduces to SEU. In the general case, a portfolio $X = (x_R, x_G, x_B)$ is evaluated as a convex combination of its minimal and its maximum expected utilities over this set Π of prior probability vectors over the three states of the world:

$$U(X) = \alpha \min_{\pi \in \Pi} \left[\sum_{s \in S} \pi(s) u(e_s\, x_s) \right] + (1-\alpha) \max_{\pi \in \Pi} \left[\sum_{s \in S} \pi(s) u(e_s\, x_s) \right] \tag{1}$$

with $\Pi = \{\pi(s) : \pi(s) \geq \underline{\pi}(s)\}$ and $s \in \{B, R, Y\}$[13]. The α coefficient can be interpreted as a measure of the agent's *aversion* to this perceived ambiguity. When $\alpha = 1$

[12]This representation of prior beliefs in the MEU model first appeared in Hey et al. (2010) and then broadly used in the ambiguity literature (see Kothiyal et al., 2014).

[13]We summarise the various sets of priors in Table 1.

the model collapses to the MEU preferences (Gilboa and Schmeidler, 1989), where maximal aversion to ambiguity is expressed. In contrast, when $\alpha = 0$, all the weight is put on the optimistic outcome. Intuitively, $\alpha > 0.5$ implies that the DM is ambiguity averse, whereas $\alpha < 0.5$ implies ambiguity seeking. Notice that in the particular framework of our study, $\alpha = 0.5$ does not imply ambiguity neutral attitudes and the model does not collapse to SEU as is the case in Ahn et al. (2014). Neutral attitudes are expressed by the uniqueness of the set Π. When this set is a singleton, the model is equivalent to the SEU and the parameter α cannot be identified.

Before presenting the different types of DMs we present how this model can be extended to its dynamic form. As is common in the ambiguity literature, this model satisfies the property of separating subjective beliefs from tastes (ambiguity attitudes). Therefore, when updating takes place, only the belief aspect of the preferences' representation is affected, while utility remains intact.

3.3 Updating Beliefs in Multiple-priors Models

We first present the updating rules for MEU, the special case of α-MEU when $\alpha = 1$. Then these rules can be naturally extended for the Hurwicz α criteria family (Hurwicz, 1951). Two ways have been suggested to update beliefs in multiple-priors models, one that satisfies DC (Epstein and Schneider, 2003; Hanany and Klibanoff, 2009; Hanany et al., 2011) and one that satisfies C (Gilboa and Schmeidler, 1993; Pires, 2002; Eichberger et al., 2007). In the former case, it suffices to solve the problem as a static one, where the allocation in the first period will determine the conditional allocation of the second period, respecting always the MEU preferences of the DM. The interesting case is when C is assumed, which allows the agent to behave in a dynamically inconsistent manner. The two most commonly updating rules include the *Maximum Likelihood Update* (MLU) and the Full Bayesian Update (FBU)[14] rule. According to the MLU rule, only the set of priors that maximise the probability of the conditional event are updated according to the Bayes rule. In the FBU rule , all the sets of priors are updated in a Bayesian way and the set of posteriors is used to evaluate the different acts. In the supplementary material, we show that in our framework with three ambiguous as-

[14] See Gilboa and Schmeidler (1993) for an axiomatisation of the rules and for references. They refer to these rules as *pseudo-Bayesian* rules.

sets, the predictions of MLU and FBU coincide. Therefore, in our analysis we assume that beliefs are updated according to the MLU rule. We now define the three non-SEU behavioural types we consider and we describe how the update rules are extended to accommodate α-MEU type preferences (when updating takes place).

3.4 The Resolute Type

The resolute type, first introduced in Hammond (1988), and later formalised in McClennen (1990) and Machina (1989) in risky contexts, embraces the simplest strategy. A resolute DM satisfies DC and the allocations at both stages coincide. This may happen for two reasons: either the DM is dedicated to somehow commit to the first stage allocations regardless the available information at $t = 1$ (aversion to information), or one can assume that beliefs are updated in a dynamically consistent manner as in Epstein and Schneider (2003). In either case, the resolute strategy with commitment is behaviourally equivalent to the dynamically consistent updating of beliefs, and to find the optimal solution, it suffices to solve the first stage problem. The optimal allocation is calculated by optimising Equation 1 subject to the budget constraint, given the DM's individual characteristics and subjective beliefs. We denote with z_s the *return* of asset s which is defined as the product between the exchange rate of the asset (e_s) and the amount of income that has been allocated to this asset (x_s): $z_s = e_s \times x_s$ with $s \in \{R, B, Y\}$. Then, in order to calculate the optimal allocation, one needs to take into consideration the relative ranking between the returns of the three assets. The various rankings depend endogenously on the amount allocated to each asset. Take for example the ranking $z_R \geq z_B \geq z_Y$[15] where red is the *best* possible outcome and yellow the *worst*. The maximum expected utility occurs at the point where the probability of the best outcome to happen is maximised (point A in Figure 1). Similarly, the minimum expected utility is obtained at the point where the probability of the best outcome R is minimised, or stating it differently, where the probability of the worst outcome Y is

[15] In total there are 7 possible rankings between the three outcomes with 6 weak inequalities and 1 strict equality.

maximised (point C). Then the α-Maxmin utility from a portfolio X is:

$$\begin{aligned} U(X) =& \alpha[\underline{\pi}(R)u(e_R x_R) + \underline{\pi}(B)u(e_B x_B) + (1 - \underline{\pi}(R) - \underline{\pi}(B))u(e_Y x_Y)] \\ &+ (1-\alpha)[(1 - \underline{\pi}(B) - \underline{\pi}(Y))u(e_R x_R) + \underline{\pi}(B)u(e_B x_B) + \underline{\pi}(Y)u(e_Y x_Y)] \end{aligned}$$

and writing the utility from the portfolio in its general form, the objective of the DM is to find an allocation X that optimises $U(X) = \sum_{s \in S} \pi(s)u(e_s x_s)$, subject to the budget and the non-negativity constraints. Here $\pi(s)$ is defined as $\pi(s) = \alpha\pi^{\min}(s) + (1-\alpha)\pi^{\max}(s)$ where $\pi^{\min}$ ($\pi^{\max}$) stands for the set of priors where the probability that the best outcome occurs is minimised (maximised). The solution of this program will provide the optimal demand for the three assets in the form $x_s^* = f(\pi, m, e, l)$, where π are now the non-additive subjective beliefs and l includes both the risk and the ambiguity attitude, which will coincide with the optimal conditional demand.

3.5 The Naïve Type

The *naïve* or *myopic* behaviour was first introduced in the literature in Strotz (1955-56), and later in Pollak (1968), indicating an agent who fails to understand the sequential nature of the problem. As a consequence, each of the stages is faced independently of the other, strategy that may lead 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 in other words, the DM solves a series of static problems and maximises his current utility. A naïve DM ignores that she is dynamically inconsistent and, as a result, the decisions that are made can potentially differ from those that had been originally planned. 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, leading to the unconditional portfolio $X = (x_R^*, x_B^*, x_Y^*)$. Then, at stage 2, she receives the partial information that one of the states did not occur, updates her prior beliefs, and based on these posteriors, she solves the maximisation problem that now involves the two remaining states, subject to the available income. Consider again the ranking $z_R \geq z_B \geq z_Y$. The DM chooses a portfolio allocation for the first period. Now assume that the partial information that the ball is not yellow ($\neg Y$) is revealed. Using the MLU rule, the DM updates those priors that maximise the probability of the event $\pi(\neg Y)$ (or $\pi(R \cup B)$). In Figure 1, this occurs in both the prior sets A and B. In addition, since the ranking of the outcomes requires that $z_R \geq z_B$[16], for the evaluation of the α-Maxmin utility it holds that $\pi^{\max} = \pi_A$ and $\pi^{\min} = \pi_B$. We denote with ${x_R}^{\neg Y}, {x_B}^{\neg Y}$ the allocations to assets R and B respectively, conditional on the information that the state is not Y. The utility of the DM inthis conditional portfolio is:

$$U(X) = \pi(R|\neg Y)u(e_R {x_R}^{\neg Y}) + \pi(B|\neg Y)u(e_B {x_B}^{\neg Y}) \tag{2}$$

with $\pi(R|\neg Y) = \alpha\pi_B(R|\neg Y) + (1-\alpha)\pi_A(R|\neg Y)$ ($\pi(B|\neg Y)$ is defined in a similar way[17]). The problem now requires us to find the conditional allocation that optimises this α-MEU, subject to both the non-negativity constraint and the new budget constraint $\hat{m}^{\neg Y} = m - x^*_Y$, where m is the initial endowed income and x^*_Y is unconditional allocation to asset Y. The conditional demand will be of the form ${x^*_s}^{\neg q} = f(\hat{\pi}, \hat{m}^{\neg Y}, e, l)$ for $s \in S \backslash q$ and $s \neq q$, where $\hat{\pi}$ is now the set of the updated beliefs.

3.6 The Sophisticated Type

Strotz (1955-56) and Pollak (1968) were among the first to recognise that pre-commitment (resolute type) is not always the optimal strategy. More specifically, the idea is that a DM who is not able to commit to her future behaviour, would prefer to adopt a strategy of *consistent planning* and then pick up the optimal plan that will actually be followed, sketching the profile of a *sophisticated* type. A sophisticated DM applies *backward induction* in order to figure out the optimal strategy for every given problem. As Hammond and Zank (2014) describe, sophistication is like the sub-game perfect Nash equilibrium of an extensive form game, between the future and present self of the DM, as in Selten (1975). Starting from the final decision nodes of a decision tree (the last period), a DM anticipates an event E to occur and therefore, the future course of action is determined by the conditional preferences of the *ex-post* self. Working backwards and applying the same principle to all the previous decision nodes, always satisfying the preferences of the *ex-post* self, she can define the optimal path that will lead her from

[16] Notice that for the naïve DM, it is not necessary for the ranking between the returns of two assets to be the same in both stages. Our estimation algorithm takes this possibility into consideration.

[17] The interested reader can consult the supplementary material where we extensively present how all the updated beliefs are calculated.

the start of the tree to the most preferable node. In this way, an optimal plan of action for the whole problem is chosen. Following this process, the DM will violate DC, as the second period optimal allocation is based on the conditional beliefs which have been updated in a dynamically inconsistent way[18]. Nevertheless, the agent is aware of this inconsistency and as Strotz (1955-56, p. 173) describes "succeeds to adopt a strategy of *consistent planning* and choose the best plan among those that he will actually follow." Siniscalchi (2011) has axiomatised this idea in the context of dynamic choice under ambiguity.

The optimal solution for the sophisticated type requires two steps. Let again the same ordering of the outcomes $z_R \geq z_B \geq z_Y$. She first solves all the three conditional problems $(\neg R, \neg B, \neg Y)$, by using her conditional beliefs and satisfying the budget and non-negativity constraints. For instance, when the information is $\neg Y$, the optimisation problem is to find the conditional allocation for assets R and B, taking into consideration the conditional beliefs, and always satisfying the outcome ranking $(z_R \geq z_B)$, and the conditional budget constraint $\hat{m}^{\neg Y} = m - x_Y^*$. The conditional allocations $x_R^{*\neg Y}$ and $x_B^{*\neg Y}$ can be written in the general form $x_s^{*\neg q} = f(\hat{\pi}, e, \hat{m}, l)$ for $s \in S \backslash q$ and $s \neq q$. Likewise, we solve for the conditional allocations for $\neg B$ and $\neg Y$. These demands are calculated in the same way as the second-stage decisions of the naïve DM (see section 3.5) and they indicate to the agent the optimal course of action for each of the conditional states (last stage of the decision problem). In the second step, the DM solves the first stage unconditional problem by taking into consideration the optimal conditional allocations, the non-negativity and budget constraint, and the relevant ranking constraint between the outcomes. The α-Maxmin utility of this two-stage portfolio is given by:

$$
\begin{aligned}
U(X) = & \frac{1}{2}\pi(\neg Y)\left[\pi(R|\neg Y)u(e_R x_R^{*\neg Y}) + \pi(B|\neg Y)u(e_B x_B^{*\neg Y})\right] \\
& + \frac{1}{2}\pi(\neg B)\left[\pi(R|\neg B)u(e_R x_R^{*\neg B}) + \pi(Y|\neg B)u(e_Y x_Y^{*\neg B})\right] \\
& + \frac{1}{2}\pi(\neg R)\left[\pi(B|\neg R)u(e_B x_B^{*\neg R}) + \pi(Y|\neg R)u(e_Y x_Y^{*\neg R})\right]
\end{aligned} \tag{3}
$$

with $\pi(R|\neg Y) = \alpha\pi_B(R|\neg Y) + (1-\alpha)\pi_A(R|\neg Y)$ and $\pi(\neg Y) = \alpha\pi_B(\neg Y) + (1-\alpha)\pi_A(\neg Y)$ where $\pi_A = \pi^{\max}$ and $\pi_B = \pi^{\min}$ are the sets of priors that satisfy the

[18] Otherwise the sophisticated strategy is identical to the resolute one.

ranking of the outcomes (the probabilities for $\neg R$ and $\neg B$ are defined in a similar way). The probability of each conditional event is multiplied by 1/2 since the subjects were informed in advance that the partial information to be revealed is randomly chosen between the two available states [19]. Notice that the conditional demands are a function of the conditional income $\hat{m}^{\neg s} = m - x_s^*$, which in turn is a function of the unconditional optimal demand for the asset s at stage 1. To calculate the unconditional demand for the three assets, it suffices to substitute the conditional income to Equation 3 and optimise with respect to the unconditional demands x_s^*.

At this point, it would be appropriate to discuss the equivalence between the α-MEU and the Chateauneuf et al. (2007) Choquet expected utility with a neo-additive capacity model (CEU). Using a notion similar to Epstein and Wang (1994) *ε-contamination* to define the set of prior distributions, Chateauneuf et al. (2007) show that the two models are equivalent in their static version[20]. This model has been empirically investigated in Dimmock et al. (2015), Baillon and Bleichrodt (2015) and Baillon et al. (2018) among others. Eichberger et al. (2010) and Eichberger et al. (2012) extend this model to its dynamic version by applying the updating rules of (Eichberger et al., 2007). Our specification is equivalent to a CEU model along with a *Generalised Bayesian Updating* rule.

3.7 Identification of Type Heterogeneity

To summarise, each of the types violates a particular axiom and this violation leads to heterogeneity in planning strategies. The resolute type satisfies DC either by updating in a dynamically consistent way or by ignoring new information but violates C. The naïve satisfies C and violates DC without realising her dynamically inconsistent behaviour. Finally, the sophisticated type satisfies C and violates DC but is aware of the potential dynamic inconsistency. Therefore, the sophisticated satisfies the consistent planning principle.

[19]This was defined by the experimental software by an independent, random draw from a uniform distribution for each of the allocation problems.

[20]In the Choquet expected utility model with a neo-additive capacity the probability of a state s is given by $P(s) = \alpha(1-\delta)\pi(s) + (1-\alpha)[(1-\delta)\pi(s) + \delta]$ where $\pi(s)$ is a reference probability distribution, α is the attitude towards ambiguity and δ is a parameter measuring the degree of confidence. In our framework the probability is given by $\pi(s) = \alpha\pi^{\min}(s) + (1-\alpha)\pi^{\max}(s)$. The models are equivalent, if one substitutes $\pi^{\min}(s)$ and $\pi^{\max}(s)$ with $\pi(s)(1-\delta)$ and $\pi(s)(1-\delta) + \delta$ respectively.

We conclude this section with a note on the identification of type heterogeneity with the use of an example. We assume a decision task with an initial endowment of $m = 100$ and for simplicity we set the exchange rates for the three assets to be equal to 1. We assume that the SEU type holds beliefs equal to 0.2, 0.5 and 0.3 for the blue, yellow and red asset respectively (for simplicity these are set equal to the objective probabilities that were used in the experiment, but any other additive distribution can be assumed). The lower probability bounds for the α-MEU types are assumed to be equal to 0.12, 0.3 and 0.18, for the blue, yellow and red asset. The lower bounds are set to be equal to 60% of the reference probability distribution of the SEU type (see footnote 20). Again, any non-additive distribution would be appropriate for this example, as long as it satisfies the ranking of the reference probability distribution. Setting the power utility coefficient equal to 0.5, we calculate the optimal allocations for all the behavioural types. We assume two levels of ambiguity attitudes: ambiguity seeking with $\alpha = 0.35$; and ambiguity averse, with $\alpha = 0.65$. The allocations are shown in figures 3a and 3b.

Both figures show the first and second-stage allocations for the blue and yellow assets, conditional on the partial information that the ball is not red ($\neg R$). The diagonal line represents all allocations where the subject can secure an equal payment from either state. The arrows in the figures show the direction of the conditional allocation. We begin with the choices of the SEU type as a benchmark. This agent is dynamically consistent, ambiguity neutral and maximises the corresponding expected utility function.

The resolute type is dynamically consistent but has ambiguity non-neutral preferences. An ambiguity seeking agent allocates more to both the worst (blue) and the best (yellow) outcome, compared to the SEU type, whilst an ambiguity averse allocates more to the worst state and less to the best. In both cases, the attitude towards ambiguity determines whether more weight is put on the worst or the best outcome. Since a resolute solves the problem as a single-stage one, the allocation remains unchanged in the second stage.

The naïve type, behaves similarly to a resolute type in the first stage, but upon the reception of the partial information, her beliefs are updated, and performs a reallocation of the available income. When the agent is ambiguity averse, the new allocation moves closer to the 45°line. This is a sign that the beliefs for the worst outcome have been updated and more weight is now put on this state compared to the unconditional stage. Therefore, the agent is willing to allocate more to this asset, behaviour which could potentially reduce her overall welfare as it was not anticipated in stage 1.

On the contrary, the sophisticated type anticipates this kind of behaviour and takes it into consideration when deciding the first-stage allocation. In both cases (ambiguity seeking and ambiguity aversion) the sophisticated agent allocates more to the worst outcome at stage 1, compared to the naïve agent, while in the second stage, the sophisticated agent allocates more to the best outcome. The intuition behind this strategy could be that the sophisticated agent takes into consideration all the possible conditional states when deciding the first-stage allocation. This agent realises that the most preferred asset (yellow) may not be available in the second stage, so she refrains from allocating the highest possible amount during the first stage. Upon reception of the partial information, the agent updates her beliefs in an anticipated way, and she allocates more to the preferred state. Indeed, the allocation of the sophisticated agent move to the opposite direction to the naïve's. Therefore, she always finds herself with higher payoffs compared to a naïve.

It becomes apparent from the Figures that a different choice pattern is predicted for each type and for each allocation problem with varying income and exchange rates. Thus, subject to asking participants a large number of allocation questions, it is possible to behaviourally distinguish between the different types. Suffice to say that for the SEU agent, all types predict the same optimal allocations and therefore are behaviourally indistinguishable. For the optimal allocations we obtain closed-form solutions by solving the corresponding maximisation problems (see Appendix).

4 Econometric Analysis

Using maximum likelihood estimation techniques, we estimate the parameters of al the specifications presented in section 3 at the subject level. This allows to introduce between and within subjects heterogeneity in three different dimensions. First, we assume heterogeneity for all the preference parameters (risk and ambiguity attitudes, beliefs, precision). Then, we allow for heterogeneity regarding the planning strategies of the agents. Finally, we allow for within-participant variation, by incorporating a random (stochastic) part in choices to capture noise in decision making.

In order to fit the various theoretical models, one needs to make several assumptions regarding the ambiguity model, the shape of the utility function and the stochastic structure of the data. Regarding the ambiguity model, we adopt the α-MEU specification for four reasons: (1) it provides a parsimonious way to capture *perceived ambiguity*; (2) the α-MEU takes into consideration both the worst and the best case scenario, providing a measure of attitude towards ambiguity; (3) well-established updating rules for the multiple-priors family of models have been axiomatised, for both the dynamically consistent and inconsistent DM; (4) kinked specifications have been shown to fit experimental data better compared to smooth ones (Ahn et al., 2014; Hey and Pace, 2014; Baillon and Bleichrodt, 2015).[21].

We turn now to the utility function. A way to obtain enough observations, and at the same time keep the structure of the decision task simple and the representation of ambiguity unchanged, is to present subjects a series of the same task with different amounts of endowment and exchange rates in each problem. A trade-off of this method is that one needs to assume a utility function over outcomes and consequently elicit the shape of this function (risk coefficient). Joint elicitation of risk and ambiguity attitudes has been employed in studies of choice under ambiguity such as Ahn et al. (2014), Hey and Pace (2014) or Baillon et al. (2018). In addition, Antoniou et al. (2015) argues in favour of controlling for risk aversion when eliciting beliefs, showing that it significantly alters inferences on deviations from Bayes Rule. Alternative designs that eliminate the use of a utility function such as pairwise choice tasks or probability matching tasks, either provide inadequate information to fit statistical models (Dominiak et al., 2012), or are not able to identify heterogeneity in planning strategies (Bleichrodt et al., 2018).

[21] We also estimated the models assuming Tversky and Kahneman (1992) Prospect Theory (Rank-dependent utility) preferences. Both models generate qualitatively similar results (similar distribution of types) but we opt for the α-MEU model as: (1) it provides an overall better in-sample fit, and; (2) there is a lack of an axiomatised updating rule for the intertemporal version of Prospect Theory.

We assume a time-invariant power utility function of the following form:

$$u(z) = \begin{cases} \frac{z^{1-r}}{1-r} & \text{if } r \neq 1 \\ \ln(z) & \text{if } r = 1 \end{cases}$$

where z is the respective payoff and r is the coefficient of risk aversion. The reasons why we favour the power form of utility are twofold. This function provides a good fit to experimental data (Wakker, 2008; Stott, 2006; Balcombe and Fraser, 2015) and also does not allow for boundary portfolios.[22].

Since the allocations are constrained to the interval $[0, m]$, a convenient way to model noise in choices is to assume that the ratios x_s/m, at a specific allocation question, are distributed according to a *Dirichlet* distribution. The random variable follows a continuous probability distribution over multinomials which are *m-tuples* $\mathbf{x} = (x_1, \cdots, x_m)$ that sum up to one. A simple parametrisation of the Dirichlet is given by setting

$$\sigma = \sum_{k=1}^{K} \beta_k$$

and

$$K = \left(\frac{\beta_1}{\sigma}, \cdots, \frac{\beta_K}{\sigma} \right)$$

with the vector K summing up to one. In our context with $K = 3$, the shape parameters of the distribution are defined as

$$\beta_s = \frac{x_s^*}{m}\sigma$$

where x_s^* is the theoretical optimal allocation to asset s. This specification has the nice property that the mean of the distribution is centered to the optimal allocation since the expression for the mean of the distribution is given by:

$$E(\frac{x_s}{m}) = \frac{\beta_s}{\sum_k \beta_k} = \frac{\frac{x_s^*}{m}\sigma}{\sigma} = \frac{x_s^*}{m}$$

where σ is the precision of the Dirichlet distribution.[23]

We then need to specify the likelihood function that will be maximised. There are

[22]Only a risk neutral or risk loving agent would choose a boundary portfolio in this particular framework.

[23]The higher the value of σ, the more precise are the choices.

two stages, with 3-way allocations in the first stage and 2-way in the second. For a particular allocation problem, let the unconditional allocation x_R, x_B, x_Y to red, blue and yellow and assume the conditional state $\neg Y$ that will lead to the conditional allocation ${x_R}^{\neg Y}, {x_B}^{\neg Y}$ and the conditional income $\hat{m}^{\neg Y}$. Using the allocations at the first stage, we assume that $x = (\frac{x_R}{m}, \frac{x_B}{m}, \frac{x_Y}{m})$ is Dirichlet distributed with the appropriate shape parameters that satisfy the properties above. The contribution to the likelihood function by the first stage allocation is given by:

$$g_1(r, \alpha, \pi, \sigma, x, x^*) = log(\Psi(\beta, \sigma))$$

where Ψ is the density function of the Dirichlet distribution,x is the vector of the actual allocations and x^* the vector of the optimal allocations.

In the second stage, there are two available allocations to be made and the dimension of the distribution is equal to 2 (it becomes a standard Beta distribution). In our example, we assume that $x^{\neg Y} = (\frac{{x_R}^{\neg Y}}{\hat{m}^{\neg Y}}, \frac{{x_B}^{\neg Y}}{\hat{m}^{\neg Y}})$ is Dirichlet distributed, subject to the suitable shape parameters. The contribution to the likelihood function by the first stage allocation is given by:

$$g_2(r, \alpha, \pi, \sigma, x^{\neg Y}, x^{\neg Y^*}) = log(\Psi(\beta, \sigma))$$

The total contribution to the likelihood function of a particular problem is given by $g = g_1 + g_2$. We consider the remaining two conditional states in a symmetric way. The likelihood function to maximise over the 60 allocation problems, is defined as:

$$\ln(\mathcal{L}(r, \alpha, \pi, \sigma, X)) = \sum_{i=1}^{60} g_i(r, \alpha, \pi, \sigma, x, x^*)) \tag{4}$$

By maximising Equation 4, we estimate the parameters using Maximum Likelihood Estimation techniques. To ensure that the solution is not trapped to a local optimum, we use a general nonlinear augmented Lagrange multiplier optimisation routine that allows for multiple restarts of the solver.[24]

We conclude this section by commenting on the number of parameters for all the

[24] The estimation was conducted using the *R* programming language for statistical computing (The *R* Manuals, version 3.0.2. Available at: http://www.r-project.org/). The estimation codes are available upon request.

specifications, as well as on the lower and upper bounds that we apply in our estimation. For the SEU specification there are four parameters to estimate, the coefficient of risk attitude r, the subjective beliefs for two out of the three states and the precision parameter σ. For the α-MEU specification, we need to estimate on top of r and σ, the set of non-additive priors π (the lower bounds) and the coefficient of ambiguity attitude α, in total six parameters. We assume either risk aversion or risk neutrality therefore, $r \geq 0$. The set of non-additive beliefs should satisfy the constraint $\underline{\pi}(R) + \underline{\pi}(B) + \underline{\pi}(Y) \leq 1$ and α is constrained to the interval [0,1], with 0 expressing extreme ambiguity seeking and 1 extreme ambiguity aversion. The maximum number of parameters to estimate is 6 and taking into consideration the amount of available choice data (120 allocation questions per subject) the obtained fit is quite stable. Finally, we assume that the type of the subjects remains stable during the experimental session and the same holds for their preferences.

5 Results

To obtain an overview of our results, we first plotted the portfolios of the subjects for each of the conditional states. Figure 2 illustrates the choices for three subjects for all the cases where the information $\neg B$ was revealed. The horizontal (vertical) axis represents the payoff if the ball is yellow (red). The 45° line stands for all the portfolio allocations that guarantee the same payoff, regardless the actual state of the world[25]. The hollow (solid) dots correspond to portfolios at period 0 (1). First, it is apparent that there is extensive violation of DC. On top of that, these violations do not seem to follow a uniform pattern, indicating the existence of heterogeneity in planning strategies.[26] The latter calls for further structural investigation.

For each subject and for each type, we have estimates of their subjective beliefs, the coefficient of risk and ambiguity attitudes (r and α), the precision parameter σ and the value of the *maximised log-likelihood*. Based on the value of the maximised log-likelihood, we can detect which type best provides the best fit to the data and therefore, classify subjects to different types. To correct for the degrees of freedom, we use both

[25] An extremely risk averse agent would always choose portfolios along this line.

[26] Indeed, subsequent econometric analysis confirmed that the left panel in Figure 2 belongs to a resolute subject (subject 13), the middle to a naïve (subject 17) and the right to a sophisticated one (subject 27).

the Bayesian Information Criterion (BIC), which controls for the different number of parameters, and the Akaike Information Criterion (AIC), which accounts for both the number of the parameters and the number of observations[27]. The average values for both measures are reported in Table 2.

In Table 3 we use the values of the maximised log-likelihood, the AIC and the BIC, to classify types at the individual level. In the first column, subjects are classified based on the fitted log-likelihood. As expected, SEU always performs worst due to the lower degrees of freedom. For the rest of the types, the sophisticated is the best for 55% of the subjects, followed by the naïve with 29% and the resolute with 16%. Columns 2 and 3 report the same information based on the corrected log-likelihoods. When AIC is used to interpret the data, the prominent type is the sophisticated (41%), followed by the SEU type (31%), the naïve (22%) and the resolute with only 5%. However, when BIC is used, 48% of the subjects are classified as SEU, followed by the sophisticated type (31%), the naïve (16%) and finally the resolute with 5%. Depending on the two information criteria, it seems impossible to make a safe inference regarding the best type that describes data, but there seems to be a consensus that the majority of the subjects are either SEU or sophisticated. Hence, we test whether the maximised log-likelihood for the best-fitting type is significantly higher compared to SEU using a *likelihood ratio test*. This test allows to compare two *nested* models where the null model is a special case of the alternative model[28]. The test statistic is given by the ratio of the two fitted likelihood functions

$$LRT = -2\ln\left(\frac{\mathcal{L}_s(\hat{\theta})}{\mathcal{L}_g(\hat{\theta})}\right)$$

whith $\mathcal{L}_s$ the maximised likelihood of the simpler model (the nested model) and $\mathcal{L}_g$ the maximised likelihood of the general model (the nesting model). The LRT statistic follows a *Chi-square* distribution with degrees of freedom $df_g - df_s$, with df_g and df_s being the number of free parameters for the nesting and the nested model respectively. With 4 parameters of the SEU and 6 of the α-MEU, the test statistic is distributed with

[27] $BIC = -2\ln(L(\hat{\theta}|x)) + k\ln(n)$, $AIC = -2\ln(L(\hat{\theta}|x)) + 2k$ 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. As is the case with the value of the log-likelihood, a lower value indicates a better fitting.

[28] Two models are nested, if the first model can be transformed into the second model by imposing constraints on the parameters of the first model. In our framework, when the beliefs are additive, the α-MEU is transformed to the SEU, so the SEU model is nested within the α-MEU.

2 degrees of freedom. Table 4 reports the classification of the subjects to types, based on the significance of the *LRT*. SEU best describes behaviour for 48% (40%) of our experimental population at 1% (5%) level of significance. For the remaining non-SEU population, the majority can be classified as sophisticated 31% (34%), naïve 16% (21%) and resolute 5% (5%).

Finding 1. *For more than half of our the experimental population, we can reject the null hypothesis of Bayesian updating at the 1% significance level. Focusing on the non-SEU subjects, the sophisticated type is best for more than half of the sample, followed by the naïve and the resolute type.*

Based on the classification of the different types, it is now possible to infer whether subjects satisfy C or DC, and when they violate DC, whether they take this into consideration or not. When we consider only the non-SEU subjects, 95% of the subjects satisfy C, while only 5% satisfy DC, which is in line with Dominiak et al. (2012).

Finding 2. *The vast majority of the experimental population with non-neutral ambiguity attitude satisfies C, while a very small percentage satisfies DC.*

We now turn to the estimates of our structural models. Table 5 reports a summary of the mean and the standard deviation of the estimated values of the parameters. We also report the median, as overfitting for few subjects may inflate the average. On aggregate, there is extensive heterogeneity regarding the values of the parameters. Figure 4 shows the distribution of the risk aversion coefficient which confirms the lack of a uniform level of risk aversion. Figures 5, 6 and 7 show the distribution of the estimated subjective probabilities for the blue, red and the yellow states respectively, for all the types (the vertical dashed line indicates the objective probability of each state). First, it seems that the distribution of the estimated beliefs when SEU is assumed, is characterised by less fat tails compared to the non-SEU types. Then, when the value of subjective beliefs is compared to the actual probabilities, it seems that subjects overestimate low probability events and under-estimate high probability events. Evidence for this finding is provided by both Table 5 and Figures 5-7. In all four cases, both the median and the average of the low probability event (B) is significantly higher compared to the actual one (0.200) which causes a right skewness to the distributions.

Likewise, the estimates for the high probability event (Y) are significantly lower compared to the objective probability (0.500) causing left skewness to the distributions. This result is in line with similar findings with a commonly observed over (under)-weighting of low (high) probability events, confirming the existence of *likelihood insensitivity*[29]. Various experiments have demonstrated the existence of this insensitivity in both student and general populations, all in static frameworks (see among others Wakker, 2010 and Abdellaoui et al., 2011). The present study, verifies the existence of this component of ambiguity attitudes, in dynamic choice frameworks.

Finding 3. *There is a systematic over-weighting of the low probability event and similarly, an under-weighting of the high probability event.*

Finally, Figure 8 illustrates the distribution of the α parameter α for the three α-MEU types. It is easy to see that there is significant heterogeneity on the attitudes, with modes at the two extremes, indicating a large number of extremely ambiguity averse people (α=1), and a considerable amount of extremely ambiguity seeking (α=0). We can then classify subjects according to their attitudes towards ambiguity. Notice that all SEU subjects are automatically classified as ambiguity neutral. For the classification of the non-SEU subjects, we consider the estimated value of α for the subject's best fitting type. 48.3% (39.7%) of the subjects are classified as ambiguity neutral, 27.6% (31%) ambiguity averse and 24.1% (29.3%) ambiguity seeking at 1% (5%) level of significance respectively. These results are in line with Charness et al. (2013), Hey and Pace (2014), Ahn et al. (2014) and Stahl (2014), all accounting for ambiguity attitudes in static frameworks. A possible explanation for the relatively high percentage of ambiguity neutral subjects is provided in Stahl (2014) where the author in footnote 2 comments on the Ahn et al. (2014) findings that "[...] one reason they may have found more expected utility maximisers that we find is that the portfolio composition task may provide context that reduces confusion among many subjects."

Finding 4. *We can reject the null hypothesis of neutral ambiguity attitudes and SEU preferences for more than half of the population. For the non-SEU agents, almost half of the population is characterised by ambiguity aversion and the remaining by ambiguity seeking attitudes.*

[29] As is explained in Trautmann and van de Kuilen (2015), likelihood insensitivity appears when people cannot distinguish between events bounded away from zero and one and transform subjective likelihoods towards fifty-fifty, resulting to an over-weighting of unlikely events and under-weighting of highly likelihood events.

Of course, one could criticise the assumption of a time-invariant utility function and hypothesise that the reason of dynamically inconsistent behaviour could be attributed to changes in the risk attitude upon reception of new information. We repeated the above analysis introducing an extra type of an expected utility maximiser with varying risk attitudes between the two stages. According to a likelihood ratio test, this model fitted best for 17% (17%) of the subjects at 1% (5%) level of significance. The proportions of the remaining types are 40% (29%) for the SEU, 24% (28%) for the sophisticated, 5% (5%) for the resolute and 14% (21%) for the naïve.

6 Conclusion

In this study we use the data from a simple two-period portfolio allocation experiment and we study heterogeneity in dynamic decision making under ambiguity. Based on the planning strategy and the axioms our subjects satisfy, we classify them to expected utility maximisers, resolute, naïve and sophisticated. Our results are summarised as: (1) almost half of the subjects behave according to the SEU model and apply Bayesian updating; (2) there is extensive violation of dynamic consistency by the non-SEU subjects; (3) the majority of the non-SEU subjects are sophisticated , few are naïve and a few are resolute, and; (4) ambiguity neutrality prevails, followed by ambiguity aversion and ambiguity seeking attitudes, almost in equal proportions.

Of course one should interpret these results with some caution. It is not possible to expect that the preference functionals presented in section 3 are the exact equations that the subjects are maximising in their minds, so the elicited behaviour should be mostly considered as an "as if" approximation. More specifically, these preference functionals, capture some crucial characteristics of behaviour that affect the way people make decisions in a dynamic environment. In particular, the SEU and the resolute type, tend to be dynamically consistent and these types are characterised by either dynamically consistent ways to update beliefs, or simply by their capacity to commit to their initial choices. The sophisticated type reflects the behaviour of the agents who anticipate that they will probably re-assess or change their preferences in the future, depending on the available information at this point, and they take this element into consideration when the make choices that include a longer horizon. Finally, the naïve decision makers are those that do not fully understand the impact that their present choices will have to their future welfare. All the types that were presented in the analysis may not fully characterise the exact decision process that each subject follows, but they seem to be successful in capturing one of the three essential elements of dynamic decision making (dynamic consistency, naïveté and sophistication).

Our results seem to provide support to those theories that assume consequentialism and abandon dynamic consistency as in Gilboa and Schmeidler (1993), Pires (2002), Eichberger et al. (2007), Eichberger et al. (2010). In particular, support is provided to theories that assume sophistication as in Siniscalchi (2011), while little support is provided in favour of theories that assume dynamic consistency (Epstein and Schneider, 2003; Klibanoff et al., 2009). The implications of the above are twofold. Recent empirical research in dynamic financial decision making, based on field data (Thimme and Völkert, 2015; Jeong et al., 2015), assumes dynamic consistency, whereas, recent theoretical studies on dynamic asset markets under ambiguity (Easley and O' Hara, 2009; Mele and Sangiorgi, 2015), assume heterogeneity in planning strategies and behaviour. Hence, accounting for heterogeneity could potentially provide better insights of how people actually behave in dynamic, ambiguous environments, fact that calls for further empirical investigation. Our paper is a first step towards studying behavioural heterogeneity regarding planning strategies in dynamic environments under ambiguity. Future research could focus on more complicated environments that include different representations of ambiguity (e.g. natural events), longer time horizons, effects of social interaction or connect it with the *decision from experience* literature, as well as with the time preferences literature.

Table 1: Prior beliefs in the MMT

π	$\pi(R)$	$\pi(B)$	$\pi(Y)$
A	$1-\underline{\pi}(B)-\underline{\pi}(Y)$	$\underline{\pi}(B)$	$\underline{\pi}(Y)$
B	$\underline{\pi}(R)$	$1-\underline{\pi}(R)-\underline{\pi}(Y)$	$\underline{\pi}(Y)$
C	$\underline{\pi}(R)$	$\underline{\pi}(B)$	$1-\underline{\pi}(R)-\underline{\pi}(B)$

Table 2: Average values of goodness of fit

Type	LL	AIC	BIC
SEU	-210.79	429.59	440.74
	(121.19)	(242.37)	(242.37)
Resolute	-208.31	428.62	445.34
	(122.14)	(244.28)	(244.28)
Naïve	-206.13	424.25	440.98
	(120.99)	(241.98)	(241.98)
Sophisticated	**-202.96**	**417.93**	**434.65**
	(119.90)	(239.81)	(239.81)
Obs	58	58	58

The Table reports the average values of the Log-likelihood, the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC) for all types. The standard deviation is reported in brackets.

Table 3: Classification based on goodness of fit

Type	LL	AIC	BIC
SEU	0	18	28
%	(0)	(0.31)	(0.48)
Resolute	9	3	3
%	(0.16)	(0.05)	(0.05)
Naïve	17	13	9
%	(0.29)	(0.22)	(0.16)
Sophisticated	**32**	**24**	**18**
%	(0.55)	(0.41)	(0.31)
Total	58	58	58

The Table reports the classification of the subjects based on the values of the Log-likelihood, the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC) for all types. The percentages are reported in brackets.

Table 4: Classification based on LRT significance

Type	Number of subjects with highest LL	Significantly different from SEU at 1%	Significantly different from SEU at 5%
SEU	0	-	-
%	(0.00)	-	-
Resolute	9	3	3
%	(0.16)	(0.05)	(0.05)
Naïve	17	9	12
%	(0.29)	(0.16)	(0.21)
Sophisticated	**32**	**18**	**20**
%	(0.55)	(0.31)	(0.34)
Non-EU	58	30	35
%	(1.00)	(0.52)	(0.60)
Total	58	58	58

The Table reports the classification of the subjects based on the Likelihood Ratio significance test (LRT), at both 1% and 5% levels of significance. The percentages are reported in brackets.

Table 5: Summary of Estimates

Parameter		SEU	Resolute	Naïve	Sophisticated
	Mean	0.279	0.254	0.228	0.222
$\underline{\pi}_B$	Median	0.298	0.281	0.261	0.260
	St.Dev	0.077	0.088	0.104	0.105
	Mean	0.370	0.346	0.318	0.337
$\underline{\pi}_Y$	Median	0.355	0.340	0.335	0.341
	St.Dev	0.060	0.059	0.100	0.111
	Mean	0.350	0.329	0.291	0.334
$\underline{\pi}_R$	Median	0.339	0.322	0.316	0.322
	St.Dev	0.059	0.059	0.099	0.073
	Mean	1.470	1.353	1.346	1.358
r	Median	0.875	0.852	0.856	0.892
	St.Dev	2.110	1.924	1.831	1.689
	Mean	-	0.478	0.482	0.502
α	Median	-	0.435	0.501	0.547
	St.Dev	-	0.417	0.407	0.370
	Mean	38.690	39.320	39.692	40.282
σ	Median	14.878	15.031	15.548	15.407
	St.Dev	52.412	52.635	52.721	52.648

The Table reports the average values, the median values and the standard deviations of the estimates for all types.

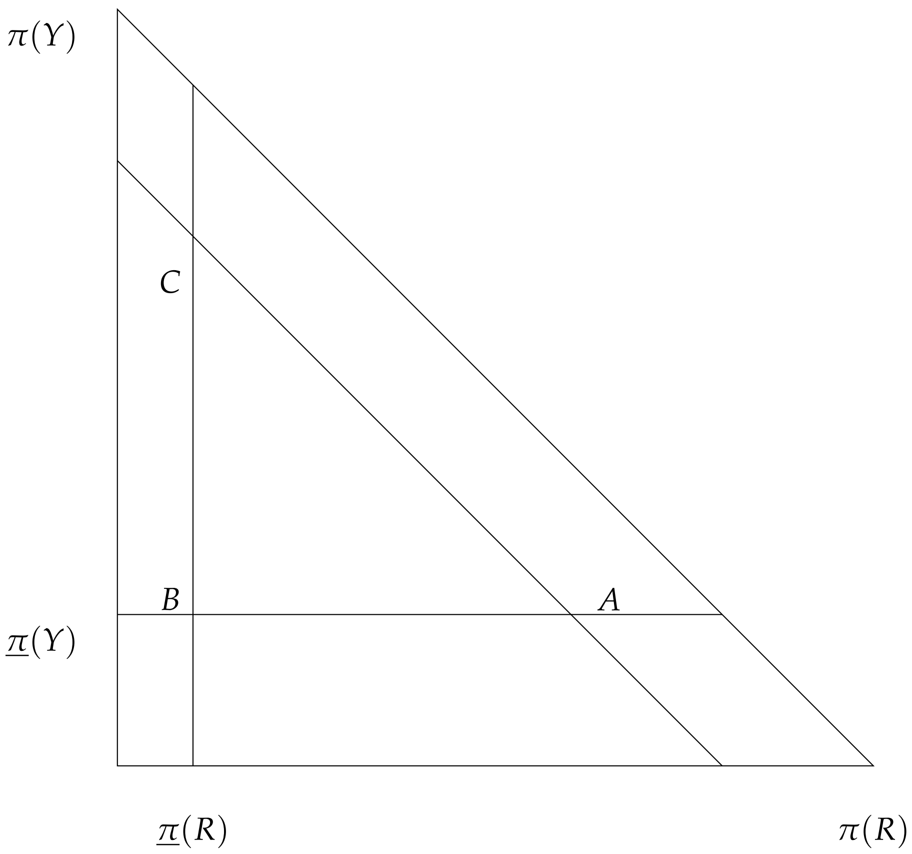

Figure 1: Prior Beliefs

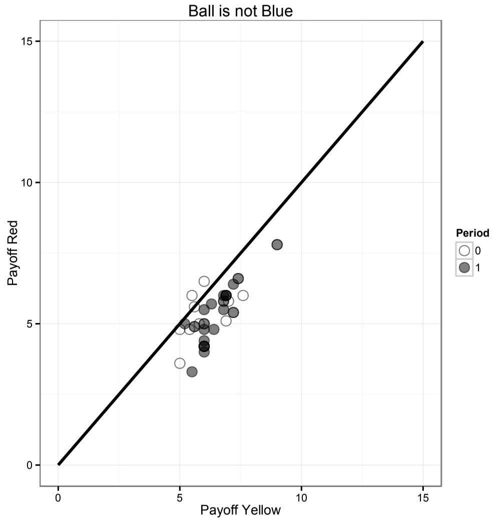

(a) Subject 13

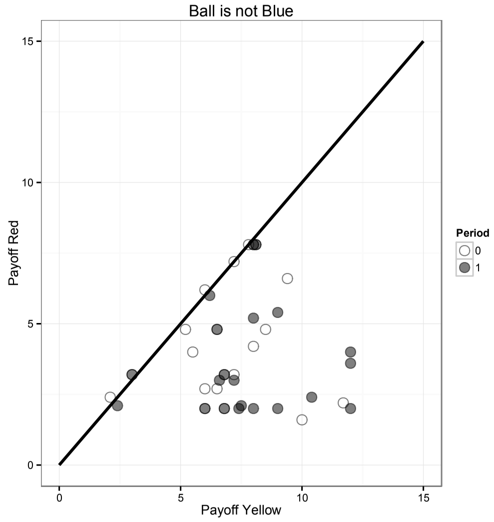

(b) Subject 17

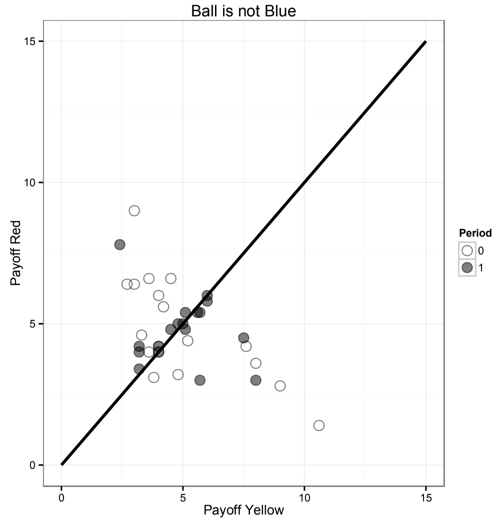

(c) Subject 27

Figure 2: Scatter-plots of Portfolios
Plotted portfolios of three representative subjects when the information $\neg B$ was revealed. The left panel shows the choices of a subject classified as resolute, the middle of a naïve and the right of a sophisticated one.

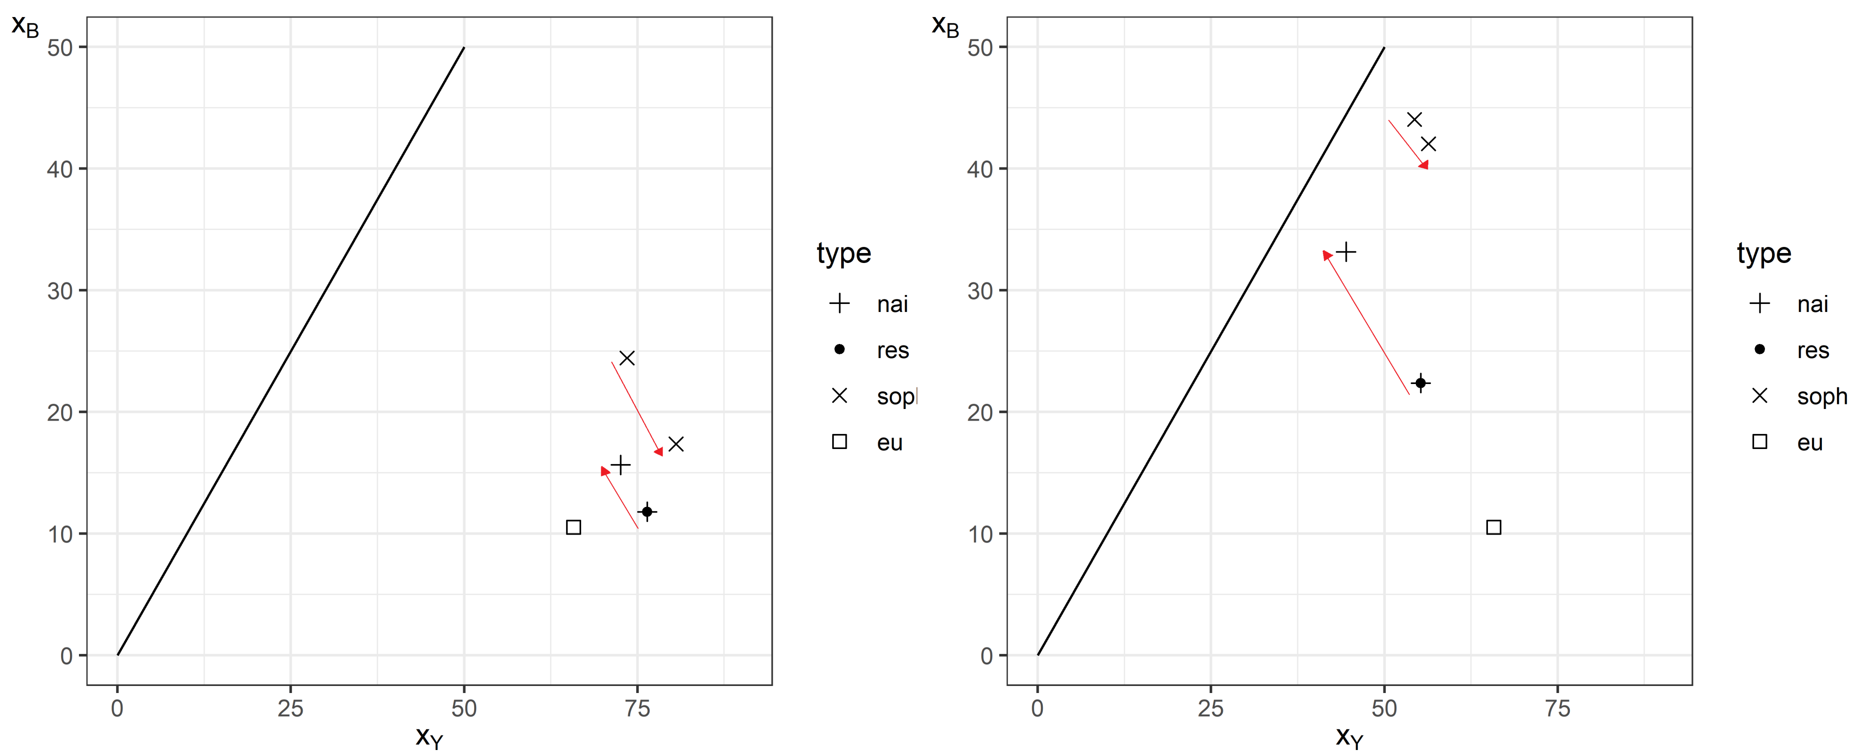

(a) $\alpha = 0.35$

(b) $\alpha = 0.65$

Figure 3: The Figures show the predicted allocations for all types in a particular task with an endowment of 100 and exchange rates equal to 1 for all the assets. Both the first and second stage allocations on the blue and yellow assets are shown, conditional on the partial information that the ball is not red. The arrows in the figures show the direction of the conditional allocation. For SEU we assume beliefs equal to 0.2, 0.5 and 0.3 for the blue, yellow and red asset respectively. For α-MEU types we assume lower probability bounds of 0.12, 0.3 and 0.18, for the blue, yellow and red asset. The power utility coefficient is set to 0.5. We assume two levels of ambiguity attitudes: ambiguity seeking with $\alpha = 0.35$ (left panel); and ambiguity averse, with $\alpha = 0.65$ (right panel).

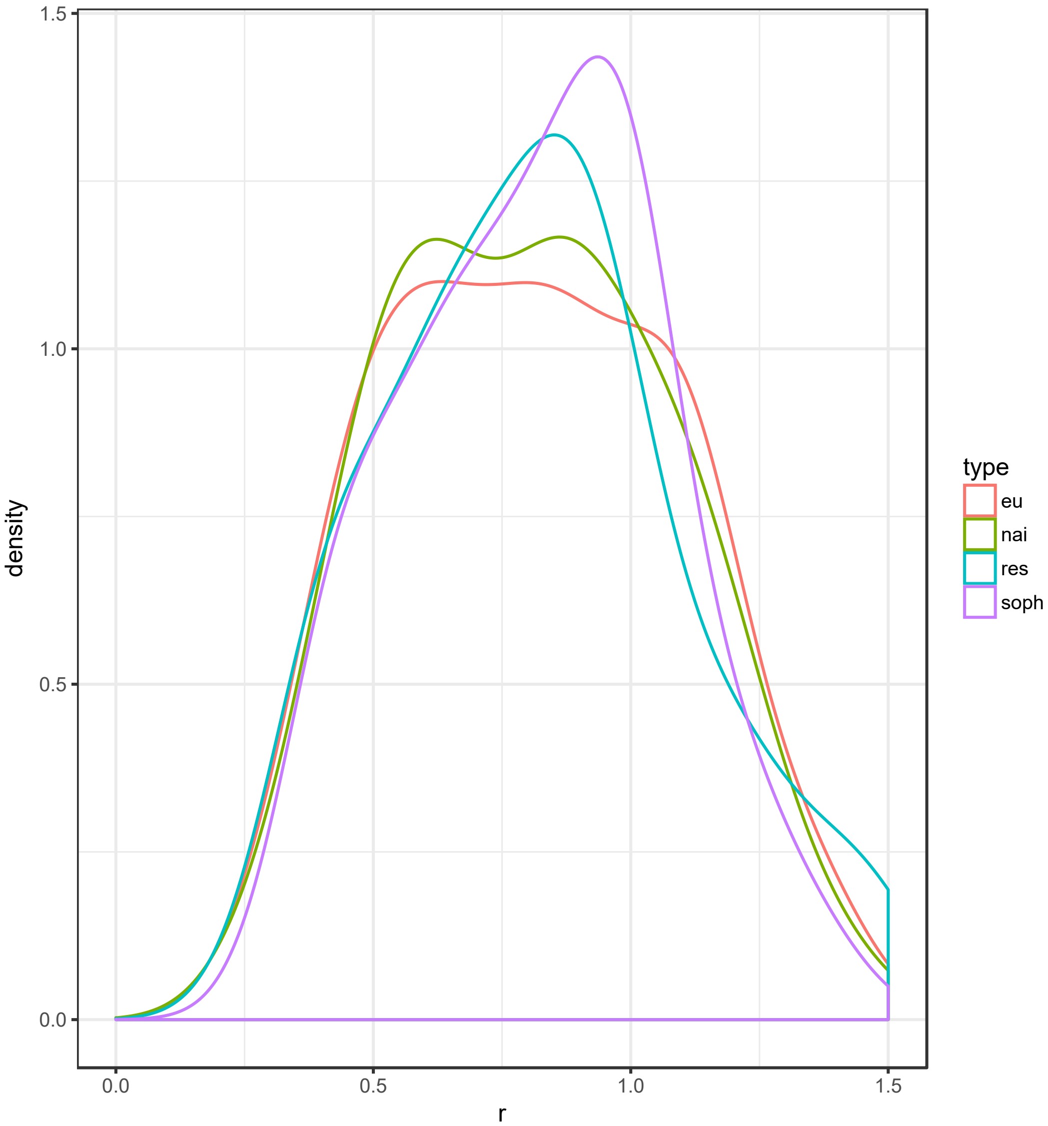

Figure 4: Kernel Distribution of the Estimated Risk Aversion Parameter

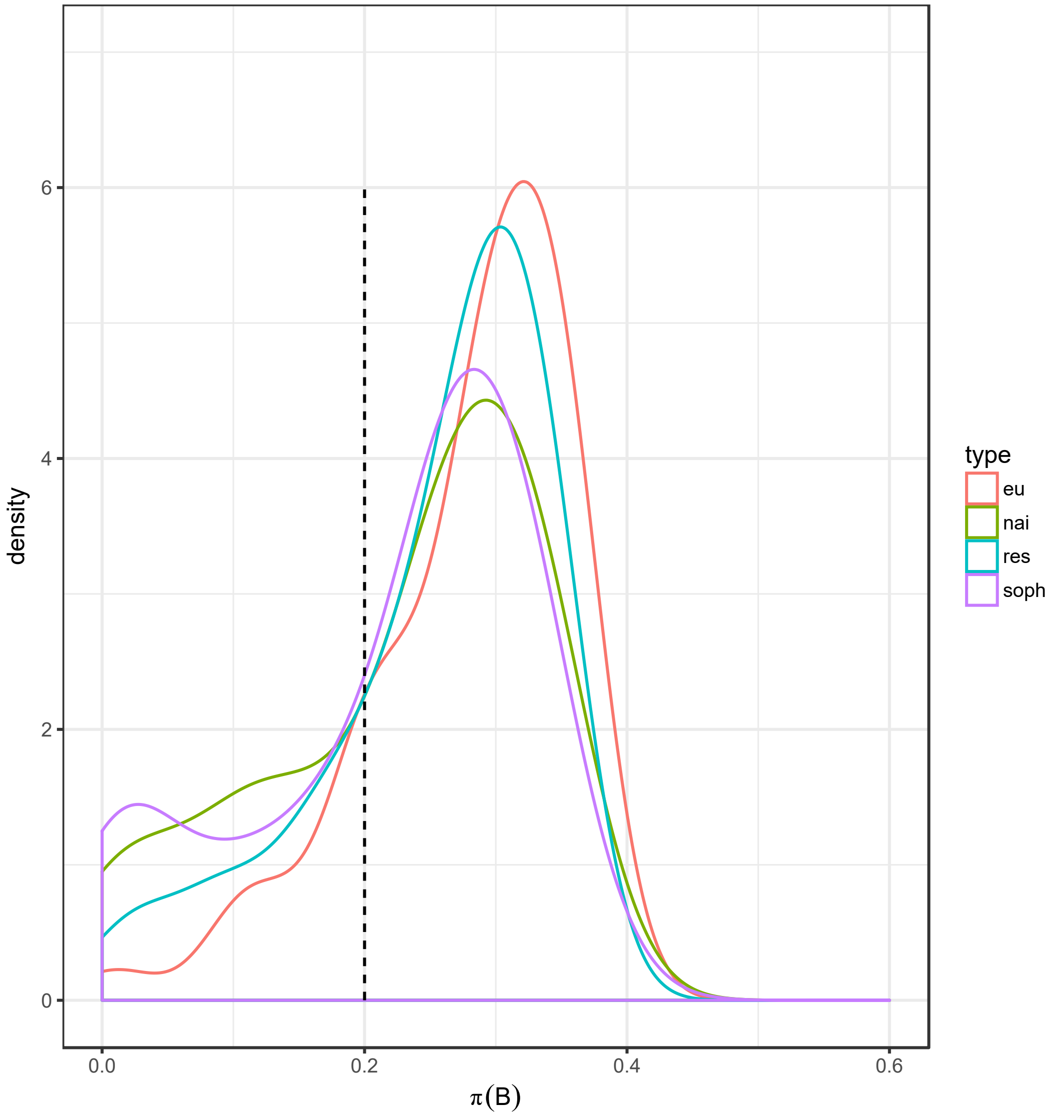

Figure 5: Kernel Density of Estimated Beliefs for the blue Asset (actual=0.2)

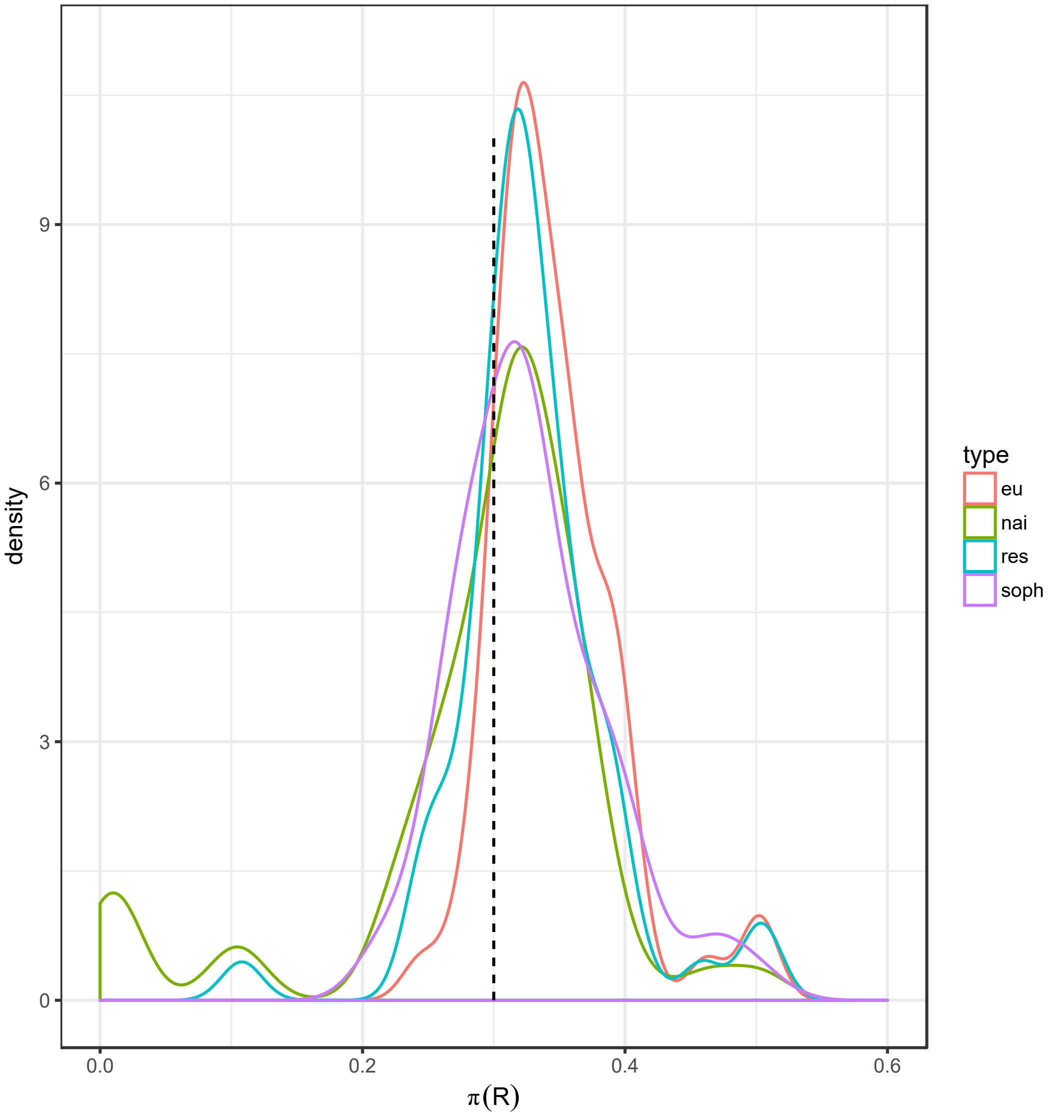

Figure 6: Kernel Density of Estimated Beliefs for the red Asset (actual=0.3)

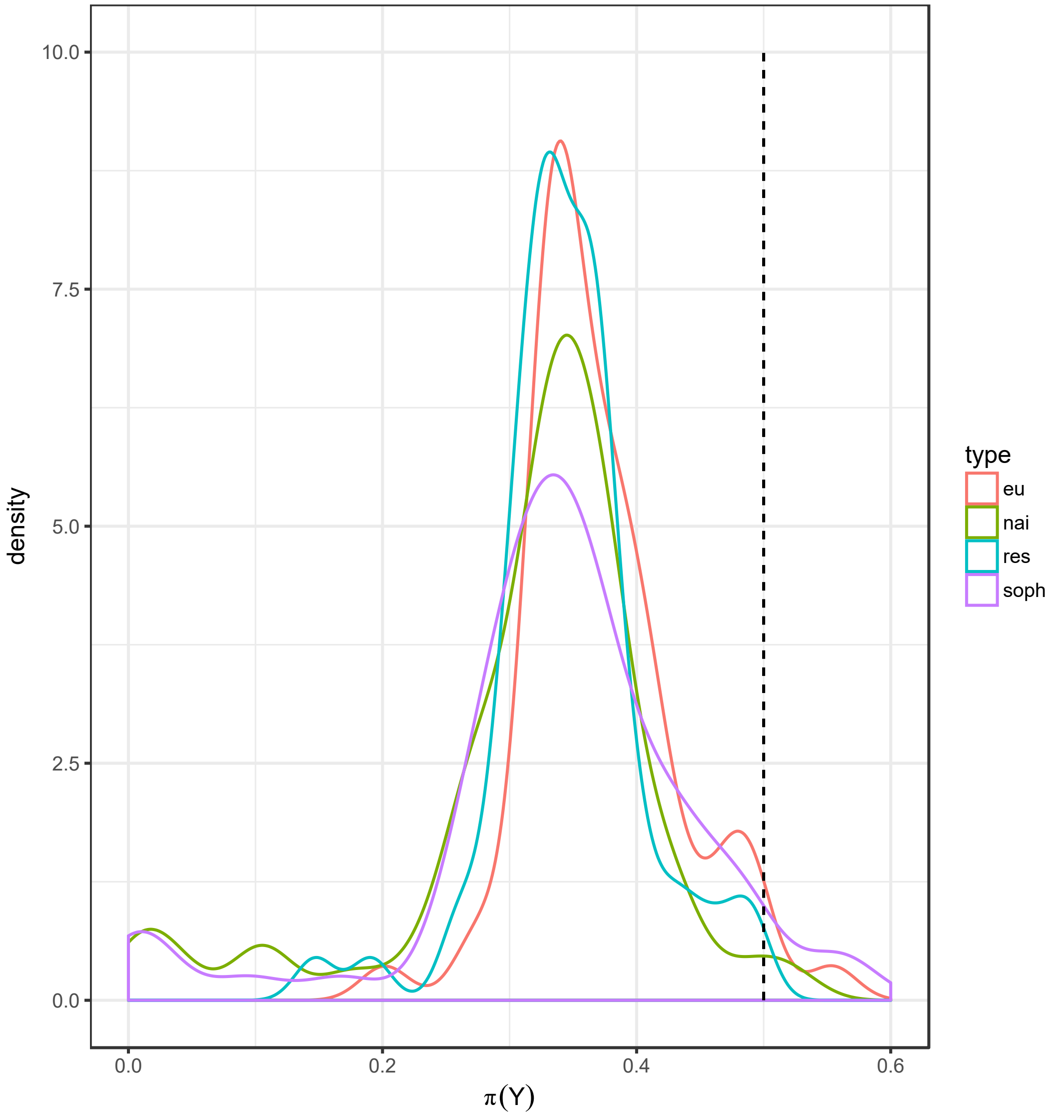

Figure 7: Kernel Density of Estimated Beliefs for the yellow Asset (actual=0.5)

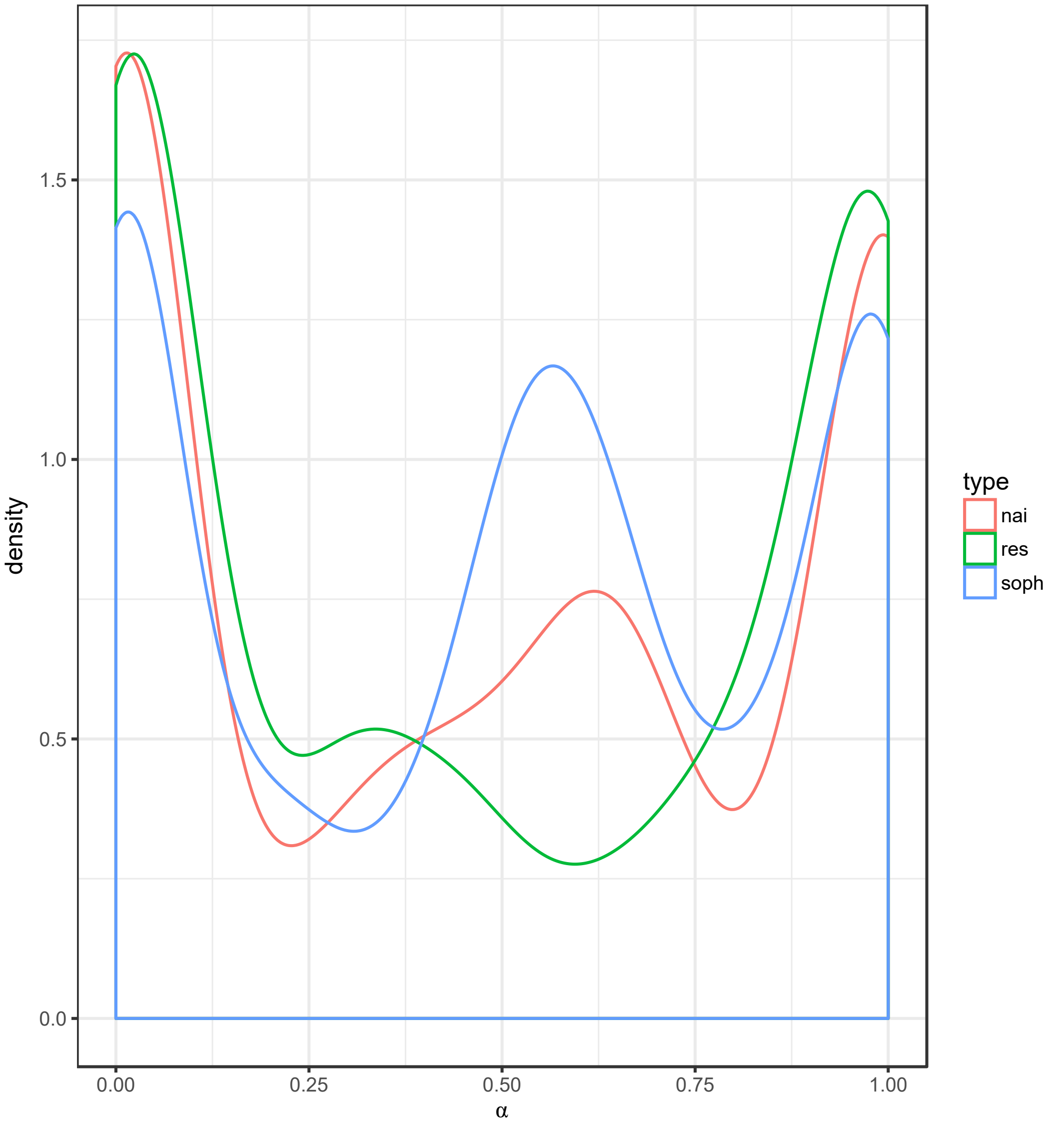

Figure 8: Kernel Density of the α parameter

References

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

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.

Antoniou, C., . Harrison, G., M. Lau, and D. Read (2015), "Subjective Bayesian Beliefs." *Journal of Risk and Uncertainty*, 50, 35–54.

Baillon, A. and H. Bleichrodt (2015), "Testing Ambiguity Models through the Measurement of Probabilities for Gains and Losses." *American Economic Journal: Microeconomics*, 7, 77–100.

Baillon, A., H Bleichrodt, U. Keskin, O. L'Haridon, and C. Li (2018), "Learning Under Ambiguity: An Experiment Using Initial Public Offerings on a Stock Market." *Management Science*, 64, 2181–2198.

Balcombe, K. and I. Fraser (2015), "Parametric Preference Functionals under Risk in the Gain Domain: A Bayesian Analysis." *Journal of Risk and Uncertainty*, 50, 161–187.

Barberis, N. (2012), "A Model of Casino Gambling." *Management Science*, 58, 35–51.

Bleichrodt, H., J. Eichberger, S. Grant, D. Kelsey, and C. Li (2018), "A Test of Dynamic Consistency and Consequentialism in the Presence of Ambiguity." Technical report.

Charness, G. and U. Gneezy (2010), "Portfolio Choice and Risk Attitudes: an Experiment." *Economic Inquiry*, 48, 133–146.

Charness, G., E. Karni, and D. Levin (2013), "Ambiguity Attitudes and Social Interactions: An Experimental Investigation." *Journal of Risk and Uncertainty*, 46, 1–25.

Chateauneuf, A., J. Eichberger, and S. Grant (2007), "Choice under Uncertainty with the Best and Worst in Mind: Neo-additive Capacities." *Journal of Economic Theory*, 137, 538 – 567.

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

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

Dimmock, S, R. Kouwenberg, O. Mitchell, and K. Peinjenburg (2015), "Estimating Ambiguity Preferences and Perceptions in Multiple Prior Models: Evidence from the Field." *Journal of Risk & Uncertainty*, 51, 219–244.

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

Easley, D. and M. O' Hara (2009), "Ambiguity and Nonparticipation: The Role of Regulation." *Review of Financial Studies*, 22, 1817–1843.

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

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 J. Lefort (2012), "Generalized neo-additive capacities and updating." *International Journal of Economic Theory*, 8, 237–257.

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

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

Epstein, L. and T. Wang (1994), "Intertemporal Asset Pricing under Knightian Uncertainty." *Econometrica*, 62, 283–322.

Ert, E. and S. Trautmann (2014), "Sampling Experience Reverses Preferences for Ambiguity." *Journal of Risk and Uncertainty*, 49, 31–42.

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

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

Georgalos, K. (2019), "An Experimental Test of the Predictive Power of Dynamic Ambiguity Models." *Journal of Risk and Uncertainty*, 59, 51–83.

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. 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.

Griffin, D. and A.. Tversky (1992), "The Weighing of Evidence and the Determinants of Confidence." *Cognitive Psychology*, 411–435.

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

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

Hammond, P. 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.

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

Hanany, E., P. Klibanoff, and E. Marom (2011), "Dynamically Consistent Updating of Multiple Prior Beliefs: An Algorithmic Approach." *International Journal of Approximate Reasoning*, 52, 1198 – 1214.

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

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*, 49, 1–29.

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

Hurwicz, L. (1951), "Some Specification Problems and Applications to Econometric Models." *Econometrica*, 19, 343–344 (abstract).

Jeong, D., H. Kim, and J. Park (2015), "Does Ambiguity Matter? Estimating Asset Pricing Models with a Multiple-priors Recursive Utility." *Journal of Financial Economics*, 115, 361 – 382.

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.

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.

Kraemer, C. and M. Weber (2004), "How Do People Take into Account Weight, Strength and Quality of Segregated vs. Aggregated Data? Experimental Evidence." *Journal of Risk and Uncertainty*, 29, 113–142.

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

Loomes, G. and G. Pogrebna (2014), "Measuring Individual Risk Attitudes when Preferences are Imprecise." *The Economic Journal*, 124, 569–593.

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

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

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

Mele, Antonio and Francesco Sangiorgi (2015), "Uncertainty, Information Acquisition, and Price Swings in Asset Markets." *The Review of Economic Studies*, 82, 1533–1567.

O'Donoghue, T. and M. Rabin (1999), "Doing It Now or Later." *American Economic Review*, 89, 103–124.

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.

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

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

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

Selten, R. (1975), "Rexamination of the Perfectness Concept for Equilibrium Points in Extensive Games." *International Journal of Game Theory*, 4, 25–55.

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.

Stahl, D. (2014), "Heterogeneity of Ambiguity Preferences." *The Review of Economics and Statistics*, 96, 609–617.

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.

Thimme, J. and C. Völkert (2015), "Ambiguity in the Cross-Section of Expected Returns: An Empirical Assessment." *Journal of Business & Economic Statistics*, 33, 418–429.

Trautmann, S. and G. van de Kuilen (2015), *Wiley Blackwell Handbook of Judgment and Decision Making*, chapter Ambiguity Attitudes, 89–116. Blackwell.

Trautmann, S. and R. Zeckhauser (2013), "Shunning Uncertainty: The Neglect of Learning Opportunities." *Games and Economic Behavior*, 79, 44 – 55.

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

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

Wakker, P. (2010), *Prospect Theory*. Cambridge University Press.

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