Comparing Behavioural Models Using Data from Experimental Centipede Games

Konstantinos Georgalos*

March 28, 2019

Abstract

The centipede game (Rosenthal, 1981) posits one of the most well-known paradoxes of backward induction in the literature of experimental game theory. Given that deviations from the unique subgame perfect Nash equilibrium generate a Pareto improvement, several theoretical models have been employed in order to rationalise this kind of behaviour in this social dilemma. The available explanations range from social preferences including fairness, altruism or cooperation motives, errors in playing, inability to perform backward induction or different depths of reasoning. In the present study, we use the Blavatskyy (2015) theoretical contribution, and relax the assumptions of Expected Utility maximisation and risk-neutral attitudes, to test an alternative explanation. We compare various probabilistic decision theory models in terms of their descriptive (in-sample) and predictive (out-of sample fit) performance, using data from experimental centipede games. We find that introducing non-Expected Utility preferences to the Quantal Response Equilibrium model, along with a non-linear utility function, provides a better explanation compared to alternative specifications such as the Level-k or the Quantal Response Equilibrium model with altruistic motives.

*Department of Economics, Lancaster University Management School, LA1 4YX, Lancaster, U.K.
✉ k.georgalos@lancaster.ac.uk, ☎ +44 (0)15245 93170
I am grateful to James Cox and Duncan James for providing me with their experimental data and to Pavlo Blavatskyy and Ivan Paya for providing helpful comments. The author would like to thank the Co-Editor of the journal Tim Salmon and a referee for very helpful comments that led to significant improvements in both the analysis and the presentation.

JEL classification: C72, C92, D81, D82

Keywords: Centipede Game, Bounded rationality, Quantal Response Equilibrium, Level-k reasoning, Probabilistic choice, Rank-dependent utility

1 Introduction

Social dilemmas in general, involve a conflict between an immediate self-interest and collective, mutual benefits. In the literature, various different frameworks have been used as a basis of social dilemmas, such as the Prisoner's Dilemma, the Chicken game, the Trust game, Public Good contributions, Common-Pool Resources games or Centipede games. All of the above frameworks share two common elements, namely the game-theoretic prediction which promotes the self-interest action as a dominant strategy, and the common findings from the experimental research in social dilemmas, that the decision makers usually deviate from the Nash-equilibrium and cooperate largely.

Krockow et al. (2016b) highlight that different explanations have been suggested to rationalise the deviations from the game-theoretic solution and classifies them to two broad categories: insufficient cognitive ability to perform backward induction reasoning, and other-regarding motives that interfere with rational, selfish decision-making. In a systematic review of centipede game experiments Krockow et al. (2016a) conclude by summarising the potential reasons why players cooperate and do not follow the rational prescription of backward induction reasoning. They include four explanations: (1) the *hidden assumption* explanation that expectations about co-players' future actions are unaffected by observations of their past actions; (2) the *cognitive burden* explanation based on the idea that backward induction is cognitively taxing and therefore players do not realise that defecting at an earlier stage may be more beneficial; (3) the *common knowledge breakdown* explanation, where players do not expect their co-players to be fully rational; and (4) the *other regarding preferences* explanation. The authors conclude that "other-regarding preferences, including prosocial behavioural dispositions and collective rationality, provide the most powerful explanation for cooperation". Van Lange et al. (2013) provide a review of the psychology of social dilemmas, and focus on recent developments in three broad areas which include the theoretical frameworks, the interdisciplinary approaches to social dilemmas, and additional factors that influence behaviour (structural, psychological and dynamic influences) in an effort to understand the theoretical puzzles and the ways cooperation is achieved.

In the present study, we focus on the developments in the theoretical frameworks as a way to explain cooperation in social dilemmas. Van Lange et al. (2013) discuss mostly theories that have been developed in the psychology literature such as the interdependence theory, the appropriateness framework and evolutionary theorising such as altruism and reciprocity. Nevertheless, the theoretical framework that has been adopted by the economic literature, in an effort to explain cooperation in social dilemmas, has mostly studied models of insufficient cognitive ability or social preferences. In a recent theoretical contribution, Blavatskyy (2015) shows, using a particular family of social dilemma games (the centipede game), that cooperation may be the result of non-Expected Utility preferences, a hypothesis that we test using data from experimental centipede games.

The centipede game, first introduced in (Rosenthal, 1981), is one of the most well known examples in the literature of social dilemmas, where the available experimental evidence significantly contradicts the predictions of backward induction. The centipede game is a dynamic game of perfect information, with a finite number of players, rounds and strategies. In its two-player version, players choose alternatively either to *Take* or to *Pass*. Playing *Take* at any decision node ends the game, while playing *Pass* gives the opportunity to the other player to choose. Each decision node corresponds to the split of a pie between the two players, with unequal payoffs, while the size of the pie increases as the players postpone to end the game for later rounds. In the particular version that we consider (see Figure 1), the game is characterised by a particular exponential payoff structure. More specifically, the pie increases exponentially (doubles) from each decision node to the next, generating social gains. Applying the cornerstone assumption of game theory, that of backward induction, along with common knowledge of rationality, predicts that all players, at all decision nodes will play *Take*, leading to the unique subgame perfect Nash equilibrium where player 1 ends the game at the very first decision node.

The centipede game induces both competitive and cooperative motives, leading to a conflict between self-interest and mutual benefits. This makes the backward induction solution look counter-intuitive, given that ending the game at a later decision node, leads to a Pareto improvement. There is a rich experimental literature confirming that subjects usually do not reason according to the backward induction logic in centipede games. A common finding is that centipede games very rarely end at the first or the last decision node, with the majority of the games ending at the intermediate nodes 4 or 5, in a six-move game. This kind of result is observed both during one-shot (or the first round of play) games, and in repeated games, where there is very little evidence of learning and convergence to Nash equilibrium (see McKelvey and Palfrey 1992). Many different explanations have been given in the literature trying to explain this finding, ranging from social preferences including fairness, altruism or cooperation motives, errors in playing, inability to perform backward induction or different depths of reasoning[1].

Additional assumptions that are usually made in game theory are that agents are Expected Utility (EU) maximisers, and that they tend to be risk neutral (have linear utility functions on payoffs). Nevertheless, empirical research in the field of decision theory, has provided strong evidence against the validity of EU, or against risk neutrality, or both. Alternative theories such at the (Quiggin, 1982) Rank Dependent Utility (RDU) or the Tversky and Kahneman (1992) Cumulative Prospect Theory might fit the data better[2]. In a recent theoretical contribution, Blavatskyy (2015) showed that introducing advances from the field of decision theory, such as non-linear probability weighting in the spirit of RDU preferences, along with a stochastic component in the choice process, may explain the deviations from backward induction.

[1]We do not aim to provide a thorough review of the literature given that there exist excellent recent reviews of the experimental literature (see Krockow et al. 2016a and Dhami 2016, chapter 12, p. 729).

[2]See Wakker (2010) and Hey (2014).

In the present study we use data from experimental centipede games to test whether alternative decision models can provide an alternative explanation of why subjects choose to continue. More particularly, using Maximum Likelihood Estimation (MLE) techniques, we econometrically explore the performance of various theoretical specifications. We investigate the performance of the models in two ways. First, we estimate models to test their in-sample performance, that is how well a model can explain a known dataset. Then, given the overfitting problems that may appear as the complexity of a model increases, we perform a cross-validation test, where we compare the predictive (out-of sample fit) capacity of each model. In addition to alternative models of choice over lotteries, we also include models of different depths of reasoning models and of social preferences. Overall, we find that relaxing the assumptions of EU preferences and risk-neutral attitudes best fits the observed deviations. The result is robust in both descriptive and predictive terms.

Our study shares similarities with McKelvey and Palfrey (1992), McKelvey and Palfrey (1995), Zauner (1999), Kawagoe and Takizawa (2012) and Ho and Su (2013) in the sense that all the aforementioned studies fit experimental centipede data to test different preference functionals. McKelvey and Palfrey (1992) provide the first experimental study in centipede games, and compare the performance of a model where every player believes that there is a small chance that their opponent is an altruist. McKelvey and Palfrey (1995) introduce the idea of the Quantal Response Equilibrium (QRE), while Zauner (1999) fits a model where the payoffs of the game are augmented by additive noise, showing an improvement against QRE and the social preferences model. Kawagoe and Takizawa (2012) test the performance of different depths of reasoning models (Level-k and Cognitive Hierarchy) showing that the Level-k model with a Poisson distribution provides a better fit compared to Cognitive Hierarchy, QRE and the social preferences model. Ho and Su (2013) compare models of different levels of reasoning (Level-k, dynamic Level-k and Cognitive Hierarchy models) in a repeated centipede game where past choices could be used to update beliefs and adapt strategies, showing that dynamic Level-k models provide a better fit. All the studies above with the exception of Ho and Su (2013), compare the models only in terms of their in-sample fit.

Our contribution to this literature can be summarised as follows: (1) using experimental data we test whether the inclusion of non-EU preferences and non-neutral risk attitudes, along with a probabilistic choice model, may provide an alternative explanation to deviations from the subgame Bayesian Nash equilibrium; (2) we compare a particular non-EU specification against the popular alternatives in the literature such as Level-k or social preferences models and; (3) we compare the models not only in terms of their in-sample fit but also in terms of their predictive capacity (out-of-sample fit).

The remainder of the paper proceeds as follows. Section 2 presents the theoretical specifications that we study. Section 3 discusses the datasets we used and the econometric methodology, while section 4 reports the results. We then conclude.

2 Theoretical Framework

In this section we present the five theoretical specifications that we investigate, the Agent Quantal Response Equilibrium model with Expected Utility preferences (EU-AQRE), the AQRE model with Quiggin (1982) Rank Dependent Utility preferences (RDU-AQRE), an extension of the RDU-AQRE where non-linear utility on monetary payoffs is assumed (RDU-AQRE-U), a behavioural model of different depths of reasoning (LEVEL-K), and an extension of the EU-AQRE where players play best responses according to their beliefs of the other player's altruistic motives (AQRE+). All the models that we compare are models of probabilistic choice, that is they consist of two parts, a deterministic part which provides a choice criterion and a stochastic part, which models the noise (errors) that a decision maker makes. The stochastic part is crucial as it facilitates econometric modelling and analysis, but it also captures an important behavioural aspect which is usually ignored in the literature. In the present study, we assume a *logit* type of errors, which is the main component of the AQRE equilibrium. Blavatskyy (2015) provides functional forms with alternative error specifications, including the Fechner model (Fechner 1860; Hey and Orme 1994), the Wilcox (2011) *contextual utility* model and the Blavatskyy (2012) probabilistic model of choice, for both EU and non-EU preferences. In our analysis we include only the specifications with a logit link function [3].

2.1 Expected Utility-AQRE

The Quantal Response Equilibrium is one of the most popular models of equilibrium behaviour in behavioural game theory, developed in McKelvey and Palfrey (1995). In this model it is assumed that best responses are not played with certainty, but players play their noisy best replies, given their beliefs about the noisy play of their opponents. This noise is captured by a logit-type error specification. The model has been extended to accommodate behaviour in extensive form games in McKelvey and Palfrey (1998) where a player is assumed to be represented by different agents, each agent deciding at a particular decision node, independently of the other agents, therefore the name Agent Quantal Response Equilibrium (AQRE).

The game is solved assuming backward induction. Using the game in Figure 1 as an example, we consider player 2 at the final decision node (node 6). Player 2 can play *Take* and get 12.80 or *Pass* and receive a payoff of 6.40. While standard game theory predicts that the player will play the best response with probability 1, AQRE predicts that due to the stochastic part of the choice process, the probability $P(T)$ of playing the dominant strategy is given by the following logistic function:

$$P(T) = \frac{\exp\left(\frac{EU(T)}{\lambda}\right)}{\exp\left(\frac{EU(T)}{\lambda}\right) + \exp\left(\frac{EU(P)}{\lambda}\right)} \tag{1}$$

with the parameter λ being a noise parameter and EU(T) the Expected Utility of *Take* (similarly

[3] We performed our analysis including all the error specifications mentioned above. To save on space we report only the results assuming a logit error specification, as this noise structure is one of the most commonly used in the literature of behavioural game theory, and it is also the one that outperformed in terms of predictive capacity. The Blavatskyy (2012) probabilistic model of choice was a specification that followed very closely. The full results are available upon request.

for EU(P)), for any given utility function $u(.)$. Notice that a high value of λ leads behaviour to random choice with $P(T) = 0.5$, while a value of λ in the neighborhood of zero predicts that the dominant strategy is played with probability very close to 1. We denote the probability that player 2 will play *Take* at node 6 as p_6.

At the penultimate node (node 5), player 1, knowing that player 2 will play *Take* with p_6, can either *Take* and receive 6.40 or play *Pass* and get 3.20 with probability p_6 and 25.60 with probability $1 - p_6$. The probability of playing *Take* in node 5 for player 1, is given by Equation 1, with $EU(T) = 6.40$ and $EU(P) = p_6 3.20 + (1 - p_6)25.60$.

Following the same methodology, it is possible to calculate the probability of each player ending the game at each decision node. Using these probabilities, one can then calculate the probability of reaching each decision node, for a given value of the noise parameter λ[4].

2.2 Rank Dependent Utility-AQRE

All the models presented above share a common characteristic, they predict that players will play *Pass* at each of the nodes with some positive probability, which is close to zero near the terminal nodes and increases as one gets towards the intermediate and initial nodes. Nevertheless, Blavatskyy (2015) shows that while these EU specifications improve upon the subgame perfect Nash equilibrium, they predict that a decision maker will never choose to play *Pass* with a probability more than 0.5 when $EU(P) < EU(T)$. However, if the decision maker overweights the probability of rare events, then the expected utility of *Pass* increases and may become greater that the utility of *Take*. A natural candidate model to accommodate this kind of behaviour is the (Quiggin, 1982) Rank Dependent Utility model [5].

The RDU model relaxes the independence axiom of EU, and allows for the probabilities to be transformed to decision weights according to a probability weighting function. Throughout

[4]The backward induction solution is presented here to illustrate how we solve the game for each of the theoretical specifications. The same analysis applies to all the remaining cases with the exception of the Level-k model, so we refrain from repeating it below.

[5]Note that since the game is defined is the gains domain the RDU models is equivalent to the Tversky and Kahneman (1992) Cumulative Prospect Theory model.

our analysis we assume a Tversky and Kahneman (1992) weighting function of the following form

$$w(p) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{1/\gamma}}$$

which transforms any objective probability p to a decision weight in the interval [0,1]. γ is a parameter to be estimated and its value defines the behavioural transformation of probabilities. When $\gamma = 1$, there is no probability distortion and the model reduces to the standard EU. When $\gamma < 1$, the probability function has an inverse-S shape and it represents the tendency of a decision maker to over (under)-weight low (high) probabilities. On the contrary, when $\gamma > 1$, the decision maker tends to under (over)-weight low (high) probabilities. In RDU, a lottery choice with two outcomes, is evaluated by ranking the payoffs from the best to the worst and calculating the expected utility after having attached the appropriate decision weights. For instance, for a lottery L which yields x with probability p and y with probability $1-p$, with $x > y$, the Rank Dependend Utility is given by $RDU(L) = w(p)u(x) + (1 - w(p))u(y)$, for any given utility function $u(.)$.

Applying this notion to the centipede example and following the backward example presented above, at node 5 in Figure 1, player 1 must choose between *Take* which yields a sure payment of 6.40, or *Pass* and get 3.20 with probability p_6 or 25.60 with probability $1 - p_6$[6]. The RDU from playing *Take* is simply $RDU(Take) = u(6.40)$, while the RDU of playing *Pass* is $RDU(Pass) = w(1-p_6)u(25.60) + (1 - w(1-p_6))u(3.20)$[7]. Since it is possible to calculate the RDU at every node, for both strategies of each player, it is possible to extend all the previous specifications and define the probability of ending the game at each decision node, for a particular set of behavioural parameters. Therefore, combining RDU with the logit stochastic error specification presented above, we obtain the RDU-AQRE specification.

[6]Remember p_6 is the probability that player 2 is playing T at node 6, which is a noisy best response defined according to one of the four specifications presented above.

[7]The same logic applies to all the previous nodes, using the compound probabilities and attaching the appropriate decision weights according to the ranking of the outcomes.

While in several experimental studies it is often convenient to assume linear utility functions on the payoffs, it is a well established result in the literature of decision making under risk, that attitudes towards risk are expressed via both the probability weighting function and the curvature of the utility function, with a concave (convex) shape of the function representing risk averse (seeking) attitudes (see Hey 2014 for a review). On top of that, there is also evidence, that failing to account for non-linearity in the utility function during the estimation of subjective beliefs, may lead to biases of the point belief estimates[8]. Goeree et al. (2003) and Goeree and Holt (2004) show that subjects exhibit moderate levels of risk aversion in games. Following Goeree et al. (2003) we adjust the models to account for risk non-neutral attitudes by assuming a constant absolute risk aversion (CRRA) function of the form $u(x) = \frac{x^{1-r}}{1-r}$ where r is a parameter to be estimated[9]. We call this specification RDU-AQRE-U.

2.3 The Level-k model

In this model, players are assumed to be heterogeneous regarding their depth of reasoning, and can be grouped to different levels (types) denoted as $L_0, L_1, L_2, \ldots L_k$[10]. The L_0 type is assumed to be a non-strategic type, that is, this type does not play any kind of best response. Every type L_k plays a best response to the belief that all the other players are of type L_{k-1}. That is, a L_1 type plays her best response to a L_0, a L_2 to a L_1 and so on.

Following Kawagoe and Takizawa (2012), in our analysis we are using a Level-k model under the assumption that the distribution of types across the population follows a Poisson distribution. More particularly, the proportion s_k of L_k type decision makers in a population is

[8] Andersen et al. (2014) show how the elicitation of subjective beliefs is improved when one accounts for both risk attitudes and probability weighting.

[9] It is a common finding in the literature that the CRRA utility provides a good fit in empirical applications (see Wakker (2008), Stott (2006).

[10] Level-k models have been applied to a variety of one-shot games. A non exhaustive list includes Stahl and Wilson (1995), Nagel (1995), Camerer et al. (2004), Costa-Gomes and Crawford (2006) and Crawford and Iriberri (2007).

determined by the following formula

$$s_k = \frac{\exp(-\tau)\tau^{k-1}/(k-1)!}{\sum_{k' \leq K} \exp(-\tau)\tau^{k'-1}/(k'-1)!}$$

with K being the total number of types. This assumption is convenient for two reasons. On the one hand, it accommodates a usually observed result from experimental studies, that the proportion of higher types is lower relative to the proportion of lower types, which can be captured by a Poisson distribution. On the other hand, this specification provides a parsimonious way to model behaviour according to a Level-k model, given that two parameters are sufficient to describe behaviour (the parameter τ for the distribution and a noise parameter λ). Furthermore, the Level-k model with a Poisson distribution appears to provide a good fit (see Kawagoe and Takizawa 2012). As a starting point, we assume that the L_0 type is an altruist type who plays *Pass* at all decision nodes. Therefore, a L_1 player 1, would play *Pass* at node 5, expecting player 2 to be L_0 and play *Pass*. Similarly, a L_1 player 2, would play *Pass* at node 4, thinking that she is playing against a L_0, and would play *Take* at node 6. Table 1 illustrates the full strategies of both players, for all levels of reasoning. We assume 6 levels of reasoning, given that the subgame perfect Nash equilibrium is achieved when player 1 is of level 6. On top of the deterministic model, we also assume that players choose their best response with some noise, which is modeled using the logistic function as above. Throughout our analysis, we denote this model as LEVEL-K.

2.4 AQRE with Altruistic Motives

The final model that we include to our analysis, is motivated by the work of McKelvey and Palfrey (1992). In this study, the authors provided an alternative explanation for their data, by transforming the game to one of incomplete information in which there is some uncertainty over the payoffs of the players. Each player in this model is assumed to believe with a very small probability ρ that their opponent is an altruist who plays *Pass* at all nodes. The model is an extension of the EU-AQRE specification presented above, with the main difference being that each player holds a prior belief ρ that the other player is an altruist. After having observed an opponent to play *Pass* in the previous node, the player updates her current beliefs regarding the altruism of the other player, and chooses accordingly. In our previous example, assuming that player 2 plays *Take* with probability p_6, at the penultimate node, player 1's belief regarding the altruism of player 2, using Bayes rule, is given by

$$\beta_5 = \frac{\rho}{\rho + (1-\rho)(1-p_4)}$$

Then, the expected utility of playing *Pass*, is given by $EU(P) = p_6(1-\beta_5)u(3.20) + (1-p_6)(1-\beta_5)u(25.60) + \beta_5 u(25.60)$. The errors in choice are assumed to follow the logistic function as above, and solving the set of nonlinear simultaneous equations for each decision node, for a given set of behavioural parameter values, we can calculating the probability of reaching each ending node. We call this specification AQRE+.

At this point it would be useful to provide an explanation of how a specification that incorporates both probabilistic choice under uncertainty and overweighting of rare events (RDU-AQRE) differs from the alternative models presented above. With the exception of the LEVEL-K model, the remaining two models (EU-AQRE and AQRE+) are based on the Expected-Utility hypothesis, augmented by a probabilistic choice rule. As Blavatskyy (2015) notes, subgame perfect equilibria derived from these models improve upon the subgame perfect Nash equilibrium. The probability of choosing P is close to zero in the last decision nodes for both players. This probability increases in the intermediate nodes but it never exceeds 0.5, at least for the 4 final nodes. As a result, the prospect P (pass) looks like the St Petersburg's lottery in the sense that it yields an ever greater payoff with an ever lower probability. As a consequence, a decision maker will never play P with a probability higher than 0.5, as long as the Expected Utility of stopping the game by playing T, is greater than playing P. Nevertheless, if a decision maker tends to overweight the likelihood of rare events, the utility of playing P, and therefore the probability of playing P at each decision node, increases (and becomes higher than 0.5 when the utility of playing P exceeds that of T). To illustrate this point, in Figure 3 we compare the predicted probabilities of a player playing P for both EU-AQRE and RDU-AQRE, at each decision node. We assume a value of λ equal to 1.6 for both models, and a value of γ equal to 0.61 for the probability weighting function (the median value from Tversky and Kahneman 1992). In both cases, the probability of playing P in the final two nodes is very close to zero. Then, it is apparent that for all the remaining nodes, the probability of playing P is always higher for the RDU-AQRE model. For instance, in node 3, the probability of P is equal to 0.36 for the EU-AQRE model, compared to 0.67 for the RDU-AQRE. This difference is attributed to the over-weighting of the low probabilities of P in nodes 4-6, which in turn increases the utility of playing P in all earlier rounds. To summarise, the RDU-AQRE predicts that a decision maker overweights the likelihood that their opponent will play P in the subsequent nodes, increasing the probability of playing P him- or herself. Consequently, this model puts larger probability mass on the intermediate nodes, compared to the alternative models, which is in line with the empirical evidence which shows that players usually stop the game at node 4 or 5.

3 Econometric Analysis

3.1 Data

Using data from experimental centipede games we compare the models presented in section 2 in two dimensions. First, we compare the models based on their descriptive validity, that is how well each model performs in terms of in-sample fitting. Then, given that overfitting may lead to wrong conclusions, we compare the various specifications based on their predictive capacity, or in other words their performance on out-of-sample fit. We use the data from the following four experimental studies: McKelvey and Palfrey (1992), Palacios-Huerta and Volij (2009), Levitt, List, and Sadoff (2011) and Kawagoe and Takizawa (2012).

The rationale behind choosing these studies lies on the fact that all share two common characteristics: (1) they use the McKelvey and Palfrey (1992) 6-move, increasing-pie centipede game with payoffs as those illustrated in Figure 1 and; (2) the game was played in its extensive form rather than its equivalent normal form, with players choosing sequentially and receiving immediate feedback of their counterpart's choice. Regarding the first characteristic, the increasing-pie game satisfies an exponentially increasing total payoff structure which generates this conflict between self-interest and social optimality and potentially provides motives to players to further *cooperate*[11]. As far as the extensive form is concerned, there are several studies where alternative methods have been used to elicit players' strategies, like asking players to indicate their preferred exit node in a reduced, normal form centipede game (Nagel and Tang, 1998), or using the *strategy* method and asking subjects in advance of their planned moves for each stage of the game as in Le Coq et al. (2015). Cox and James (2012) use two variations of the game adopting an independent private values (IPV) information environment. They represent the game by employing two different institutional formats (clock or tree) and two different dynamic structures (sequential or simultaneous move). While each variation provides the flexibility to obtain more data per subject, they also carry additional disadvantages as it is not straightforward whether subjects perceive the various forms of the game as equivalent[12]. In addition, with the exception of McKelvey and Palfrey (1992), in all the other studies the game was played only once, without providing the subjects the opportunity to learn and adapt their strategies. This is in line with the spirit of Costa-Gomes and Crawford (2006) who claim that modelling initial responses in games is of paramount importance in providing insights into cognition and strategic behaviour. All of the aforementioned studies, share a common experimental design, that is players are randomly matched in pairs without knowing the identity of the other player, players make choices sequentially in the game with the payoffs illustrated in Figure 1[13] and all choices were financially incentivised. The set of data we analyse consists

[11]Note that in a constant sum centipede game, the size of the pie does not alter throughout the game and playing *Take* at node 1 is the only strategy that leads to a fair allocation.

[12]See Brants and Charness (2011) for an exposition on this topic.

[13]Most recent studies have scaled up the payoffs of McKelvey and Palfrey (1992) by 100. This does not seem to

of both lab and field experiments. McKelvey and Palfrey (1992) and Kawagoe and Takizawa (2012) conducted traditional lab experiments, Palacios-Huerta and Volij (2009) provide data from both a field study with professional chess players and lab sessions with chess players and student participants, while Levitt, List, and Sadoff (2011) analyse data from a field experiment with chess players from two international open chess tournaments. This gives us in total five independent datasets (three lab and two field studies). The raw data from each experimental study are shown in Table 5. [14]

3.2 In-sample fit

To compare the in-sample fit performance of the models, we estimate all the theoretical specifications using the pooled dataset from the four studies. This gives us a total of 295 observations (590 subjects). In the pooled set, it is apparent that the subgame perfect Nash equilibrium assumption can be safely rejected, with only 3% of the pairs ending the game at node 1, while more than 50% ending the game at nodes 4 and 5. Figure 2 illustrates this result. We estimate our models using standard Maximum Likelihood Estimation (MLE) techniques. More particularly, for each of the models, we are able to work out the probability $p_i(v)$ of playing *Take* at each decision node i according to the assumptions of each theoretical specification regarding both the deterministic and the stochastic part of the decision making process, given a particular set v of behavioural parameters. The latter provides a vector $P(v)$ of length n for the n-move game containing all $p_i(v)$ for $i = 1, 2, \ldots, n$. Then, we are able to calculate the probability $q_i(P(v))$ of reaching each decision node i (e.g. $q_1 = p_1(v), q_2 = (1 - p_1(v))p_2(v)$ and so on). Following Zauner (1999) and Kawagoe and Takizawa (2012), the log-likelihood function for each model can be written as

$$\mathcal{L}(v) = \sum_{i=1}^{n+1} n_i \ln(q_i(P(v))) \tag{2}$$

affect our econometric analysis and therefore we are able to pool all the data together.

[14]The data have been retrieved from the Appendix in Kawagoe and Takizawa (2012). As in Kawagoe and Takizawa (2012), we also exclude the field experiment of Palacios-Huerta and Volij (2009) with chess players, given the remarkably high level of backward induction play.

where n_i is the observed frequency of playing *Take* at decision node i. Using MLE techniques, we are able to estimate the parameter vector $\hat{v}$ that maximises Equation 2. Then, based on the value of the maximised likelihood, along with appropriate goodness of fit measures and significance tests, we can select models based on their in-sample fit performance[15]. For simplicity, we assume homogeneity in the parameters for both players.

3.3 Out-of-sample fit

Comparing models based only on their descriptive adequacy (in-sample fit) is susceptible to drive to misleading conclusions. More specifically, more complex models are prone to *over-fitting*. That is, while the dataset may be generated by a simple model plus some noise, a more complex model is able to fit this noise and provide an artifactually better fit. Over-fitting is a particular characteristic of analysis conducted with both lab and field data due to their limited sample size and the fact that tests based on asymptotic distributions are not valid for these small sample size datasets. An alternative model selection method is to use the *K-fold cross-validation* (rotation estimation) method (Geisser, 1993)[16]. The cross validation technique is able to assess how well the results of an econometric estimation will generalise to an independent dataset. More concretely, this method evaluates the accuracy with which a model can predict choices in an unknown dataset. To this end, a fold is formed by partitioning the whole sample to two subsets, the *training* set (70-80% of the sample) which is used to estimate parameters for the models, and the *validation* set (20-30% of the sample), which is excluded from the estimation and is used to assess the predictive accuracy. The same process is repeated K times, for each of the K folds, where in each fold a different part of the data is used as the training (validation)

[15] 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/). To avoid local optima, a multiple-restarts routine was applied, with different starting values at each restart. The routine is using a general non-linear optimization using the augmented Lagrange multiplier method. (package *rsolnp*, Ghalanos and Theussl (2012)). The estimation code is available upon request.

[16] For applications on model selection and validation see Zhang (1993) and Keane and Wolpin (2007). For a recent discussion on overfitting and predictive accuracy see Wilcox (2007) and Stahl (2018). Recent studies of model selection based on the predictive power of the models include Hey et al. (2010) and Kothiyal et al. (2014).

test, resulting to the calculation of an appropriate statistical measure. Note that we partition data at the level of datasets, the entire dataset from a particular experimental study appears either in the training set or the validation set, but not in both. The validation set is always an entire dataset from an experimental study, that has not been included in the training set.

In our context, the training set of each fold is used to estimate the vector $\hat{v}$ of behavioural parameters. This set $\hat{v}$ is then applied to the validation dataset to generate the implied *Take* probability vector. Then, using this probability vector along with the actual *Take* frequencies of the validation set, we are able to calculate the out-of-sample predictive log-likelihood $\mathcal{L}_{OUT}(\hat{v})$. The model selection criterion instructs to choose the model with the highest predictive log-likelihood. Given that we have five independent datasets, we form the 5 folds, each consisting of four pooled datasets as the training set, and the one dataset left out, as the validation set. We rotate the datasets in such a way that each dataset is used for validation exactly once and as a training set K-1 times (Table 6 summarises the data partition). We compare models based on the aggregated value of their out-of-sample predictive log-likelihood value

$$\sum_{k=1}^{K} \mathcal{L}_{k,OUT}(\hat{v})$$

3.4 Identification of models

Early work on comparing theories of learning in games (Salmon 2001; Wilcox 2006) has shown that a particular family of learning models posses difficulties in identifying the actual data generating process of simulated datasets. On the other hand, similar studies that compare models of choice under uncertainty (Hey and Pace 2014; Kothiyal et al. 2014), have been more successful in showing that it is actually possible to identify the underlying true model, making this kind of comparisons more meaningful. To test whether the models are properly identified and are indeed behaviourally distinguishable, we conduct an extensive Monte Carlo simulation, similar to the one in Wilcox (2006) and Hey and Pace (2014)[17]. For each of the models, we

[17] Analytical details on the simulation are provided in the online Appendix.

assume a reasonable set of parameters and we generate a simulated dataset (predicted probabilities of playing T), which is then used to individually estimate each of the 5 specifications. We use the Akaike Information Criterion (AIC) as a model selection criterion[18]. It is expected that, if identification is possible, whenever the model assumed in the estimation coincides with the data generating model (true model), the lowest AIC value will be generated, compared to the AIC value that all the other specifications generate (best fit of the data corrected for the degrees of freedom). We repeat this process for a number of 100 simulations. The results confirm that the data-generating model is always identified, and the underlying parameters are successfully recovered.

3.5 Empirical content of the AQRE models

Another possible criticism relates to whether the estimated Quantal Response Equilibrium models carry any empirical content. In a recent paper, Haile et al. (2008) show that QRE imposes no falsifiable restrictions, that is, it can rationalise any distribution of behaviour in any normal form game. Nevertheless, under the particular structural assumptions we make in the parameterisation of the models, there are three potential arguments against this criticism. First, from a theoretical point of view, in order for the result in Haile et al. (2008) to be valid (Theorem 1, page 184), one should relax the assumption of i.i.d. perturbations across each player's strategies. In our analysis, we maintain the standard assumption that the experimental literature is making on i.i.d. perturbations from a logit parametric family, when we specify and estimate the AQRE models (as in McKelvey and Palfrey 1992; Kawagoe and Takizawa 2012 and Wright and Leyton-Brown 2017). Then, based on our simulation exercise, if the AQRE models are able to rationalise any set of probabilities, and therefore are not falsifiable, one would expect that they would provide a fit to the data, at least as good as any of the other competing models, independently of the actual data-generating process. The results from the

[18]The AIC is defined as $2k - 2\ln(L)$, with k the number of parameters and $\ln(L)$ the value of the maximised Log-likelihood. A lower value of AIC indicates a better fit.

simulations do not provide any evidence in favour of this. On the contrary, the parameter recovery exercise, shows that it is possible to both identify the AQRE models, and to recover the corresponding distribution value of the parameters. Finally, our estimation/prediction methodology shows that AQRE models carry empirical content. If the AQRE models carry no empirical content, they will perform pretty well in the in-sample fitting (potentially due to over-fitting) but because of this over-generalisation they will have a very poor predicting performance. The results reported in the next section, show that those specification that perform well in the in-sample analysis, perform well in the out-of-sample prediction, implying that the empirical content that AQRE models carry is not due to over-fitting, and therefore verifying their usefulness in empirical work.

4 Results

The results from our econometric analysis are presented in two parts. First, we rank models based on their in-sample performance and we discuss model selection using the value of the maximised log-likelihood, along with goodness of fit measures as criteria. In the second part we present the results of the K-fold validation exercise, where models are ranked according to their predictive capacity.

All the results from the estimations are presented in Table 2. The top panel of the Table reports the implied probabilities of reaching each terminal node. All of the models allocate some positive probability mass to each of the outcomes, with the main characteristic being the unimodal distribution of the outcomes. In particular, most of the models predict a very low probability of the game ending in one of the two extreme outcomes (play *Take* at 1 or *Pass* at 6) while all of the models put most of the probability mass to nodes 4 and 5, which is in line with the actual observed frequencies as is illustrated in the second column of the Table (play *Take* at nodes 4 or 5). The middle panel, reports the MLE estimates of the parameters for each model along with their standard errors. The bottom panel includes measures of goodness of fit. The first row reports the value of the maximised loglikelihood. A higher value indicates a better fit of the model. Nevertheless, given that the models under consideration differ in their degrees of freedom, we also report the values of three information criteria which correct for the different degrees of freedom and the size of the sample, namely the *Akaike Information Criterion* (AIC), the *corrected Akaike Information Criterion* (AIC_c) and the *Bayesian Information Criterion* (BIC).[19] A lower value of each of the criteria indicates a better fit of the model.

While EU-AQRE predicts a very high probability of the game ending at node 1 (0.256), RDU-AQRE-U and LEVEL-K attach a lower probability to this outcome and correctly assign most of the probability mass to outcomes 4 and 5. As a consequence, the two latter models provide a better fit to the data, fact which is reflected by the lower values of the AIC (1042.1 and 1045.1). Based on the AIC, the RDU-AQRE-U model seems to provide an overall best fit. Nevertheless, one may ask whether this difference is statistically significant or not. In all subsequent analysis we use two types of significance tests, the *Vuong* test (Vuong, 1989) for non-nested models and the *Likelihood ratio* test for nested models[20].

We start by testing whether the assumption on non-linear probability weighting improves the in-sample fit of the EU-AQRE specification. Using a Likelihood ratio test within models (EU-AQRE Vs. RDU-AQRE), with one degree of freedom, it is confirmed that RDU performs significantly better than EU (p<0.000 at the 1% significance level). Observing the estimated value of the probability weighting parameter γ, the parameter appears to be significantly dif-

[19]The standard AIC is defined as $2k - 2\ln(L)$, with k the number of parameters and $\ln(L)$ the value of the maximised Log-likelihood. While this measure corrects for the degrees of freedom, it does not take into consideration the sample size, leading often to overfitting. (Burnham and Anderson (2004)). Sugiura (1978) and Hurvich and Tsai (1989) proposed a modified version of AIC that also corrects for the sample size (second-order bias). The AIC_c is defined as

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

with n being the sample size. Finally, the BIC is defined as $2k\ln(n) - 2\ln(L)$.

[20]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. EU-AQRE nests all the other AQRE specifications, while LEVEL-K and AQRE are non-nested models.

ferent than 1 ($\gamma = 1$ indicates linear probability weighting and therefore EU preferences). The value of γ (0.714) implies the theoretical intuitive shape of the weighting function (for $\gamma < 1$ the function has an inverse S shape, indicating over-weighting of low probability events).

Relaxing the assumption of risk neutrality and assuming a concave utility function on monetary payoffs, along with non-linear probability weighting (RDU-AQRE-U) leads to a dramatic improvement in the in-sample performance of the model. The value of the AIC drops from 1126.5 to 1042.1, and a Likelihood ratio test confirms that this difference is statistically significant (p<0.000), showing that the curvature of the utility function may explain an important part of the observed choices. The estimate of the risk curvature parameter r, is significantly different than 0 and provides evidence of moderate risk aversion, with r being equal to 0.866. Focusing on the probability weighting parameter γ, the estimate is in line with the intuition (and the common evidence from experimental studies) of an inverse S-shaped weighting function ($\gamma = 0.744$). Figure 4 illustrates the shape of the weighting function and the different weighting of low and high likelihood events for both RDU-AQRE and RDU-AQRE-U.

Up to now, the analysis shows that the RDU model with non-linear utility performs significantly better compared to the EU specification. We now turn to the question, how well does this RDU specification perform, compared to the popular alternatives from the literature of behavioural game theory. Kawagoe and Takizawa (2012) compare the standard AQRE model (under the assumption of EU along with risk neutrality) against the Level-k family of models and the altruistic version of the AQRE (AQRE+). The two last columns of Table 2 report the estimates of the LEVEL-K[21] and the AQRE+ specifications. AQRE+ is ranked third over all the specifications, showing that results probably are not driven by the players' beliefs on their opponent's altruistic motives (the estimated parameter ρ is equal to 0.033), while the LEVEL-K

[21]The estimated parameter τ of the Poisson parameter shows the following distribution (percentage) of levels in the population $L_1 = 16.7, L_2 = 29.9, L_3 = 26.7, L_4 = 15.9, L_5 = 7.1, L_6 = 2.6$, confirming the often observed result of the majority being classified as levels 1-3.

model is ranked second and very close to the best fitting model RDU-AQRE-U. This difference is statistically significant at the 5% level according to a Vuong test (p=0.049).

When in-sample fit performance is used as a model selection criterion, the results show that a RDU specification can explain the data better, closely followed by the LEVEL-K model. Nonetheless, one could argue that this difference could be attributed to the over-fitting, caused by the increased complexity of the model. To complete the analysis, we compare the models using their predictive capacity as a criterion. Following the K-fold validation process, Table 3 reports the predicted log-likelihood for all the specifications that we have encountered in this study. The Table reports the value for each of the five folds in the first five columns, with the two last columns reporting the summed predictive log-likelihood and the mean value (cross-validated likelihood) respectively. As a reminder, a higher value of the predictive likelihood (a less negative value) indicates a stronger predictive capacity[22]. The results show that the best predictive model is the RDU-AQRE-U specification with a cross-validated likelihood of -95.7 and an aggregated value of -478.4. with the LEVEL-K being classified second with a cross-validated likelihood of -97.7 and an aggregated value of -488.4, followed by the AQRE+. EU-AQRE, has the worst predictive capacity.

As a robustness test of the out-of-sample fit, we use the results obtained above and perform a final out-of-sample test on data that have not been yet used in the analysis. More specifically, we use the estimated values of the models to calculate the predicted likelihood in two experimental datasets of 4-move centipede games. The 4-move centipede game is identical to the one in Figure 1, with the only difference that ends at node 4 where players may end with payoffs of 6.40 and 1.60 respectively, if they both play "Pass" at all nodes. We use the data from the McKelvey and Palfrey (1992) 4-move game, as well as data from Treatment 1 in Cox and James (2015) who replicate McKelvey and Palfrey (1992). The results are reported in Table

[22] As this is the predictive log-likelihood, one does not need to take into consideration different degrees of freedom.

4 and confirm the conclusions of the analysis above. RDU-AQRE-U and LEVEL-K are the two specifications which outperform all other models, while the EU-AQRE model is ranked last in both cases. Nevertheless, one should be careful when interpreting these results, given that both datasets include a relatively small number of observations (10 and 20 pairs respectively) and almost all the games end in either the second or the third node, with zero mass in the three remaining outcomes. The latter point, argues in favour of the *K-fold cross-validation* as a model selection method, which identifies a winner model, that can on *average* predict behaviour in a better way compared to the competing models.

5 Discussion and Concluding Remarks

In this study, we use data from experimental centipede games to test the idea of Blavatskyy (2015) that non-EU preferences may provide an alternative explanation of why subjects choose *Pass* in later decision nodes rather than at the first node, as backward induction instructs. Introducing a stochastic part to the decision making process to capture errors in choice, along with a deterministic model, we form probabilistic decision making models. We then compare the various specifications in two ways, first in terms of their descriptive capacity (in-sample fit), that is how well a model can explain behaviour in a known dataset, and then in terms of their predictive capacity (out-of-sample fit), that is how well a model can predict behaviour in an unknown dataset. In addition to these specifications, we fit two of the most common models in the literature that have provided an explanation of behaviour in centipede games, namely the LEVEL-K model and a version of the Quantal Response Equilibrium with altruistic motives.

The results show that a non-EU specification along with a concave (risk averse) utility function provide an alternative explanation to the exit frequencies in centipede games. More particularly, the RDU -AQRE-U, the specification which includes RDU preferences along with a logit stochastic error and a power utility function (CRRA) outperforms all other models, both in terms of in-sample and out-of-sample performance. The second best model is the LEVEL-K model, a popular model in the behavioural game theory literature. Both outperform the EU-AQRE model that has also been extensively used in modelling behaviour.

These results seem to resemble the results from the field of dynamic decision making under risk, where the existence of non-EU preferences leads to dynamic inconsistencies, fact which leads to violations of backward induction (see Barberis 2012; Hey and Panaccione 2011). Our analysis shows that relaxing the standard assumptions of game theory, namely Expected Utility maximisation along with risk-neutrality, may provide a better explanation to the available data, compared to limited cognitive ability or social preferences models. More particularly, introducing the assumption of non-linear probability weighting, allows to model the behavioural aspect of overweighting the likelihood of rare events, providing an alternative explanation of why players keep passing in the centipede game, when they should not.

As Eichberger et al. (2018) highlight, "[$\cdots$] there have been attempts to modify Nash equilibrium by introducing random deviations (Quantal Response Equilibria, k-level equilibria) in order to obtain a better fit for experimental data. Though, more flexible due to the extra parameters, these concepts provide no interpretation for these parameters which could allow one to make ex-ante predictions". Blavatskyy (2015) shows that the probabilistic choice models could be benefited by introducing overweighting of the likelihood of rear events. Our analysis shows that this modification, not only increases the in-sample fit, which would be expected as the models becomes more flexible, but it also enhances the predictive power of the model. Indeed, it appears that overweighting carries an actual economic meaning, given that in all the specifications, the predictive capacity of the model dramatically increases when non-linear probability weighting is allowed.

The LEVEL-K model differs from the AQRE specifications, in the sense that it first makes exact predictions of what players of different levels will play, and then estimates a distribution of the different types found in the dataset. As a result, it will always perform well in the in-sample fit. Nevertheless, the only information one receives regarding behavioural traits of the subjects, is the percentage of types in a population, information which is not straightforward how it can be used to help understanding behaviour in different games and environments. On the contrary, there is a plethora of converging empirical evidence regarding the degree of risk aversion and probability weighting of experimental subjects, in various situations. It also worths mentioning that the best fitting specification (RDU-AQRE-U) is one of the most commonly used in analysing experiments of individual choice under risk, and is proven to provide a good fit to the data. Combining the two points above, could make the results of the present study portable to further applications.

The present study shows that introducing advances from the experimental literature of individual decision making, provides choice models with improved explanatory and predictive power for an important class of social dilemma games, the centipede games. It is a first step towards providing support in favour of an alternative theoretical explanation of why people tend to cooperate in social dilemmas, showing that it is possible to construct a cooperative equilibrium, based only on "selfish" behaviour. These results could contribute to our better understanding of how people engage in strategic reasoning. First, adopting a probabilistic decision making model which allows for non-linear probability weighting may be able to better explain behaviour in centipede games that employ a different format or institution (Nagel and Tang 1998; Le Coq et al. 2015; Cox and James 2012). Then, these results could be used to explain behaviour in games that share the same structure with the centipede game (i.e. sequential two-player games with complete and perfect information). Examples of the latter include the Rubinstein (1982) sequential bargaining game, the Selten (1978) chain store paradox or the Kreps and Wilson (1982) repeated prisoner's dilemma. We leave this question for future work.

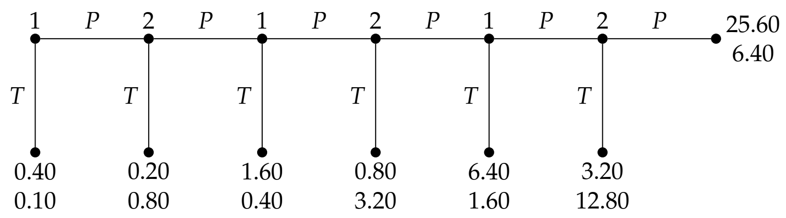

Figure 1: The 6-move centipede game McKelvey and Palfrey (1992).

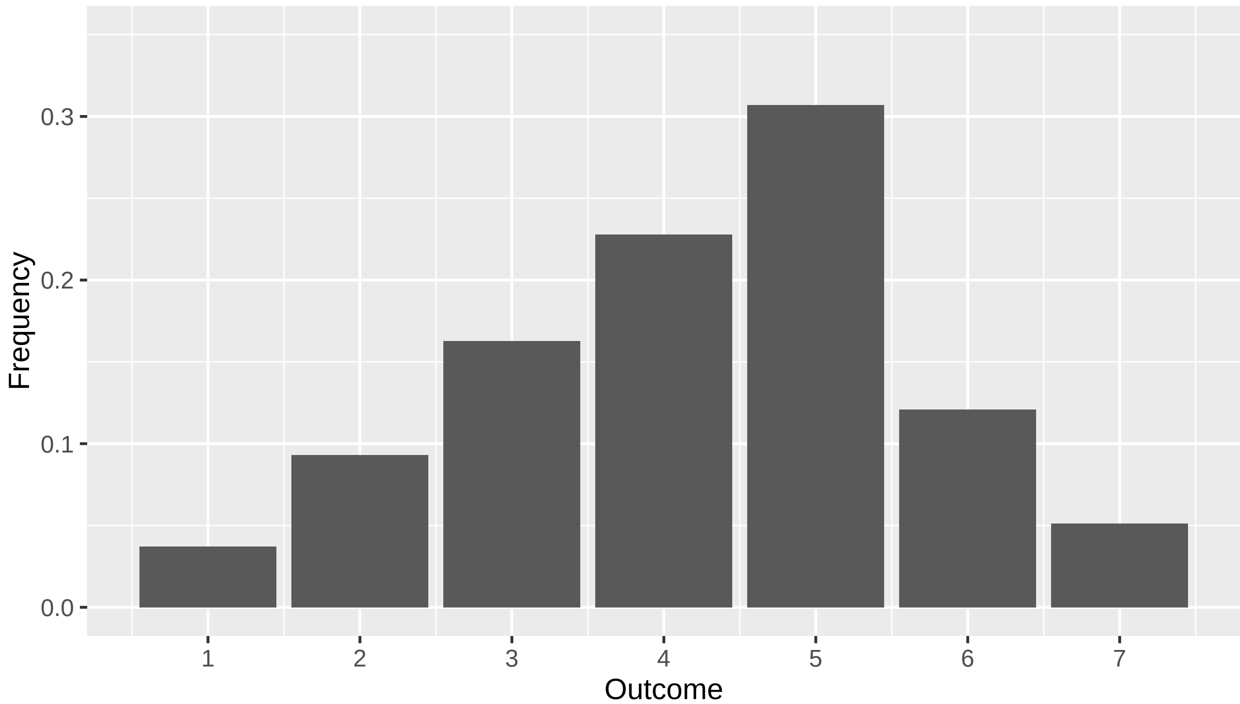

Figure 2: Distribution of playing *Take* at each decision node in the six-move centipede game in the experimental data used in this study.

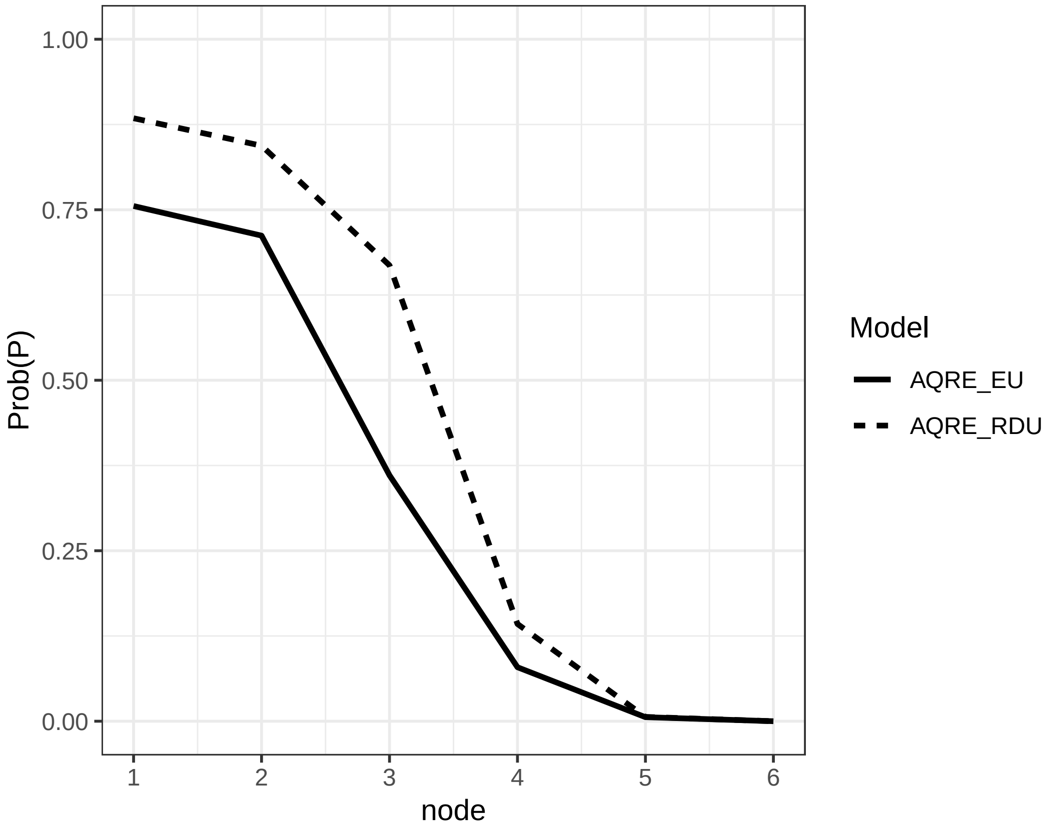

Figure 3: Probability of playing *Pass* (P) at every decision node.

Level	Player 1	Player 2
L_0	P,-,P,-,P,-	-,P,-,P,-,P
L_1	P,-,P,-,P,-	-,P,-,P,-,T
L_2	P,-,P,-,T,-	-,P,-,P,-,T
L_3	P,-,P,-,T,-	-,P,-,T,-,T
L_4	P,-,T,-,T,-	-,P,-,T,-,T
L_5	P,-,T,-,T,-	-,T,-,T,-,T
L_6	T,-,T,-,T,-	-,T,-,T,-,T

Table 1: Rule hierarchy for the LEVEL-K model in the 6-move centipede game. Player 1 makes decisions at nodes 1, 3 and 5 {1,-,3,-,5,- } while player 2 at nodes 2, 4 and 6 {.,2,-,4,-,6}. P and T stand for the strategies *Pass* and *Take* respectively. In this model each player of Level n plays her best response against a player of Level $n-1$. For example, a player of Level 3 will play T at node 5 as she thinks she is playing against a Level 2 player, who plans to play T at node 6.

Outcome	Dist	EU-AQRE	RDU-AQRE	RDU-AQRE-U	LEVEL-K	AQRE+
1	0.033	0.256	0.145	0.060	0.026	0.127
2	0.119	0.183	0.116	0.127	0.094	0.125
3	0.214	0.134	0.128	0.214	0.203	0.111
4	0.237	0.149	0.161	0.265	0.289	0.173
5	0.271	0.203	0.331	0.220	0.22	0.339
6	0.088	0.071	0.116	0.099	0.078	0.122
7	0.037	0.004	0.003	0.015	0.09	0.003
λ		2.131	1.719	0.066	0.018	1.712
s.e.		0.105	0.109	0.039	0.268	0.101
τ		-	-	-	1.789	-
s.e.		-	-	-	0.071	-
ρ		-	-	-	-	0.033
s.e.		-	-	-	-	0.055
γ		-	0.714	0.744	-	-
s.e.		-	0.043	0.053	-	-
r		-	-	0.866		
s.e.		-	-	0.068		
LL		-591.348	-561.232	-518.037	-520.529	-560.282
AIC		1184.695	1126.463	1042.074	1045.058	1124.564
AIC_c		1184.709	1126.504	1042.156	1045.099	1124.605
BIC		1188.382	1133.837	1053.135	1052.432	1131.938
Obs		295	295	295	295	295

Table 2: In-sample comparison for all the specifications. The top panel of the Table reports the implied probabilities of reaching each terminal node, the panel in the middle reports the parameter estimates and the bottom panel the goodness of fit measures.

Model	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5	Overall	Mean
EU-AQRE	-57.668	-91.917	-217.170	-67.150	-133.357	-567.262	-113.452
RDU-AQRE	-51.982	-74.199	-202.609	-65.440	-122.615	-516.845	-103.369
RDU-AQRE-U	-46.9972	-67.5331	-184.764	-61.4585	-117.613	-478.365	-95.673
LEVEL-K	-46.836	-62.675	-184.267	-66.451	-128.155	-488.384	-97.677
AQRE+	-50.169	-63.271	-203.584	-63.644	-126.471	-507.140	-101.428

Table 3: Comparison of the predictive performance of all the models.

Model	MP92-6	CJ
EU-AQRE	-17.246	-19.176
RDU-AQRE	-16.73	-18.777
RDU-AQRE-U	-13.044	-14.665
LEVEL-K	-12.78	-14.794
AQRE+	-16.985	-18.785

Table 4: Out-of-sample performance of the models in the 4-move centipede game. The abbreviations correspond to the data from the original study: MP92-4 McKelvey and Palfrey (1992), CJ Cox and James (2015).

Outcome	MP92-6	KT	LLS	PHV-F	PHV-L	ALL
1	0	1	4	3	2	10
2	3	1	10	6	15	35
3	3	1	17	14	28	63
4	9	3	25	12	21	70
5	10	25	27	4	14	80
6	3	10	12	1	0	26
7	1	3	7	0	0	11
Total	29	44	102	40	80	295

Table 5: Raw data from all the experimental studies used in the analysis.

Fold	Training set	Validation set
1	KT, LLS, PHV-F, PHV-L	MP92-6
2	MP92-6, LLS, PHV-F, PHV-L	KT
3	MP92-6, KT, PHV-F, PHV-L	LLS
4	MP92-6, KT, LLS, PHV-L	PHV-F
5	MP92-6, KT, LLS, PHV-F	PHV-L

Table 6: Folds in the fold-K cross validation. The abbreviations correspond to the data from the original study, MP92-6 McKelvey and Palfrey (1992), PHV Palacios-Huerta and Volij (2009) F for the field and L for the lab study, LLS Levitt, List, and Sadoff (2011) and KT Kawagoe and Takizawa (2012).

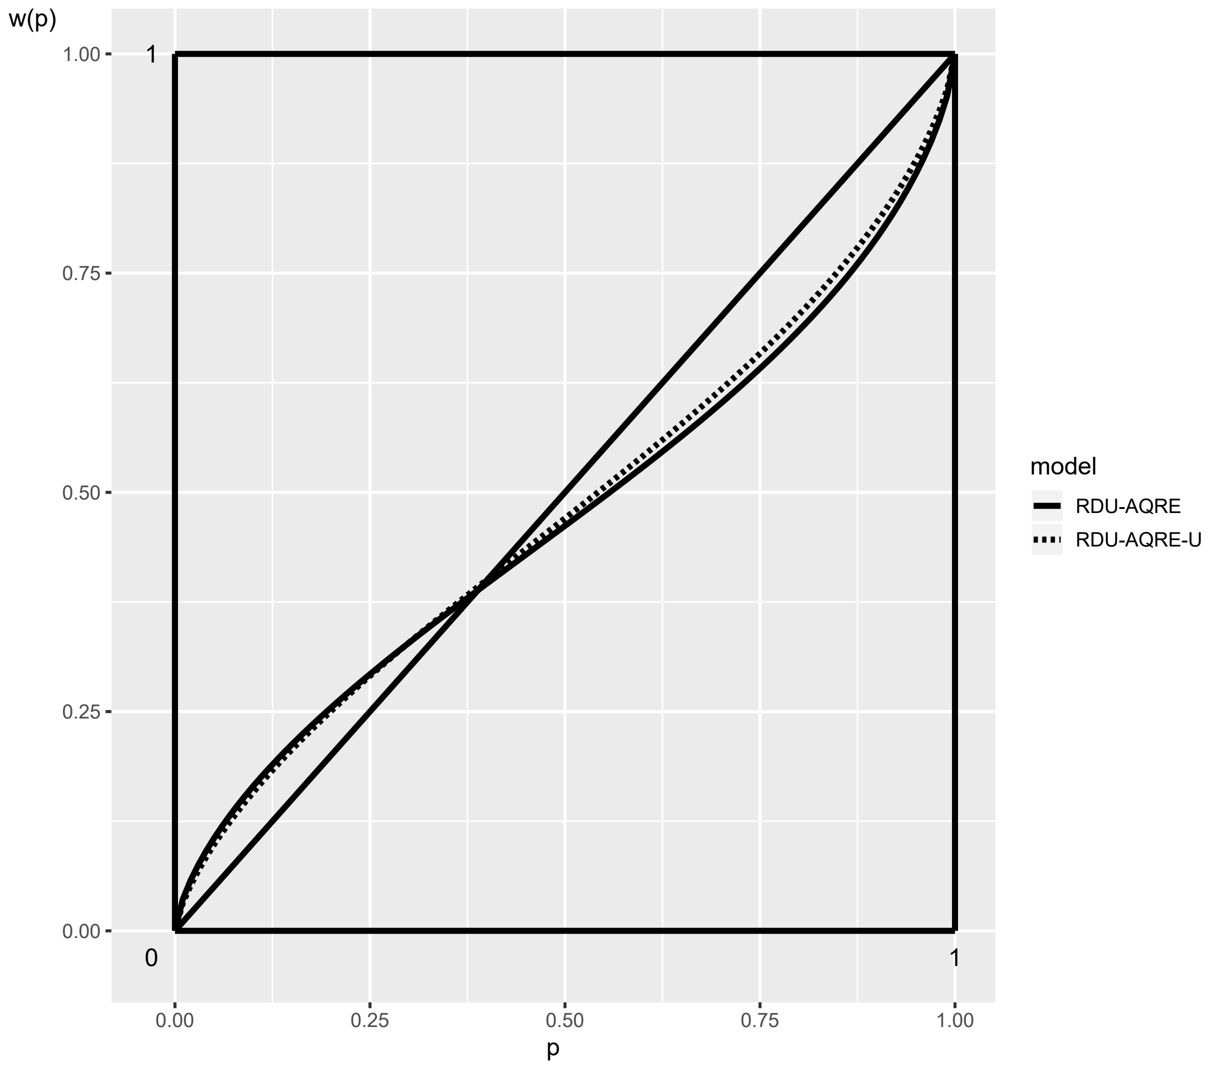

Figure 4: Probability weighting functions with (RDU-AQRE) and without (RDU-AQRE-U) linear utility.

References

Andersen, S., J. Fountain, G. Harrison, and E. Rutström (2014), "Estimating Subjective Probabilities." *Journal of Risk and Uncertainty*, 48, 207–229.

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

Blavatskyy (2012), "Probabilistic Subjective Expected Utility." *Journal of Mathematical Economics*, 48, 47 – 50.

Blavatskyy, P. (2015), "Behavior in the Centipede Game: A Decision-theoretical Perspective." *Economics Letters*, 133, 117 – 122.

Brants, J. and G. Charness (2011), "The Strategy Versus the Direct Method: A First Survey of the Experimental Comparisons." *Experimental Economics*, 14, 375–398.

Burnham, K. and D. Anderson (2004), "Multimodel Inference Understanding AIC and BIC in Model Selection." *Sociological Methods & Research*, 2, 261–304.

Camerer, C., T-H. Ho, and Chong J. (2004), "A Cognitive Hierarchy Model of Games." *Quarterly Journal of Economics*, 119, 861–898.

Costa-Gomes, M. and V. Crawford (2006), "Cognition and Behavior in Two-person Guessing Games: An Experimental Study." *American Economic Review*, 96, 1737–1768.

Cox, J. and D. James (2012), "Clocks and Trees: Isomorphic Dutch Auctions and Centipede Games." *Econometrica*, 80, 883–903.

Cox, J. and D. James (2015), "On Replication and Perturbation of the McKelvey and Palfrey Centipede Game Experiment." In *Replication in Economic Experiments* (G. Deck, E. Fatas, and T. Rosenblat, eds.), volume 18, 53–94, Emerald, Research in Experimental Economics.

Crawford, V. and N Iriberri (2007), "Fatal Attraction: Salience, Naivete, and Sophistication in Experimental Hide-and-Seek Games." *American Economic Review*, 97, 1731–1750.

Dhami, S. (2016), *The Foundations of Behavioral Economic Analysis*. Oxford University Press.

Eichberger, J., S. Grant, and D. Kelsey (2018), "Ambiguity and the Centipede Game: Strategic

Uncertainty in Multi-Stage Games with Almost Perfect Information."

Fechner, G. (1860), *Elements of Psychophysics*. Holt, Rinehart and Winston.

Geisser, S. (1993), *Predictive Inference*. New York, NY: Chapman and Hall.

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

Goeree, J. and C. Holt (2004), "A Model of Noisy Introspection." *Games and Economic Behavior*, 46, 365 – 382.

Goeree, J., C. Holt, and T. Palfrey (2003), "Risk Averse Behavior in Generalized Matching Pennies Games." *Games and Economic Behavior*, 45, 97 – 113.

Haile, P., A. Hortaçsu, and G. Kosenok (2008), "On the Empirical Content of Quantal Response Equilibrium." *American Economic Review*, 98, 180–200.

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., 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 C. Orme (1994), "Investigating Generalizations of Expected Utility Theory Using Experimental Data." *Econometrica*, 62, 1291–1326.

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.

Ho, T. and X. Su (2013), "A Dynamic Level-k Model in Sequential Games." *Management Science*, 59, 452–469.

Hurvich, C. and C. Tsai (1989), "Regression and Time Series Model Selection in Small Samples." *Biometrika*, 297–307.

Kawagoe, T. and H. Takizawa (2012), "Level-k Analysis of Experimental Centipede Games." *Journal of Economic Behavior & Organization*, 82, 548–566.

Keane, M. and K. Wolpin (2007), "Exploring the Usefulness of a nonrandom Holdout Sample for Model Validation: Welfare Effects on Female Behavior." *International Economic Review*, 48, 1351–1378.

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.

Kreps, D. and R. Wilson (1982), "Reputation and Imperfect Information." *Journal of Economic Theory*, 27, 253–279.

Krockow, E., A. Colman, and B. Pulford (2016a), "Cooperation in Repeated Interactions: A Systematic Review of Centipede Game Experiments, 1992-2016." *European Review of Social Psychology*, 27, 231–282.

Krockow, E., A. Colman, and B. Pulford (2016b), "Exploring Cooperation and Competition in the Centipede Game through Verbal Protocol Analysis." *European Journal of Social Psychology*, 46, 746–761.

Le Coq, C., J. Tremewan, and K. Wagner, A. (2015), "On the Effects of Group Identity in Strategic Environments." *European Economic Review*, 76, 239–252.

Levitt, S.D., J.H. List, and S.E. Sadoff (2011), "Checkmate: Exploring Backward Induction Among Chess Players." *American Economic Review*, 101, 975–990.

McKelvey, R.D. and T.R. Palfrey (1992), "An Experimental Study of the Centipede Game." *Econometrica*, 60, 803–836.

McKelvey, R.D. and T.R. Palfrey (1995), "Quantal Response Equilibria for Normal-form Games." *Games and Economic Behavior*, 10, 6–38.

McKelvey, R.D. and T.R. Palfrey (1998), "Quantal Response Equilibria for Extensive-form

Games." *Experimental Economics*, 1, 9–41.

Nagel, R. (1995), "Unravelling in Guessing Games: an Experimental Study." *American Economic Review*, 85, 1313–1326.

Nagel, R. and F. Tang, F. (1998), "Experimental Results on the Centipede Game in Normal-form: an Investigationation on Learning." *Journal of Mathematical Psychology*, 42, 356–384.

Palacios-Huerta, I. and O. Volij (2009), "Field Centipedes." *American Economic Review*, 99, 1619–1635.

Quiggin, J. (1982), "A Theory of Anticipated Utility." *Journal of Economic Behavior and Organization*, 3, 323–343.

Rosenthal, R. (1981), "Games of Perfect Information, Predatory Pricing and the Chain Store Paradox." *Journal of Economic Theory*, 25, 92–100.

Rubinstein, A. (1982), "Perfect Equilibrium in a Bargaining Model." *Econometrica*, 50, 97–109.

Salmon, T., C. (2001), "An Evaluation of Econometric Models of Adaptive Learning." *Econometrica*, 69, 1597–1628.

Selten, R. (1978), "The Chain Store Paradox." *Theory & Decision*, 9, 127–159.

Stahl, D. (2018), "An Empirical Evaluation of the Toolbox Model of Lottery Choices." *The Review of Economics and Statistics*, 100, 528–534.

Stahl, D. and P. Wilson (1995), "On Players' Models of Other Players-Theory and Experimental Evidence." *Ganes and Economic Behaviour*, 10, 218Ű254.

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

Sugiura, N. (1978), "Further Analysis of the Data by Akaike's Information Criterion and the Finite Corrections." *Communications in Statistics, Theory and Methods*, 13–26.

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

Van Lange, P., J. Joireman, C. Parks, and C. Van Dijl (2013), "The Psychology of Social Dilemmas: a Review." *Organizational and Human Decision Processes*, 120, 125–141.

Vuong, H., Q. (1989), "Likelihood Ratio Tests for Model Selection and non-nested Hypothesis." *Econometrica*, 57, 307–33.

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.

Wilcox, N. (2006), "Theories of Learning in Games and Heterogeneity Bias." *Econometrica*, 74, 1271–1292.

Wilcox, N. (2007), "Predicting Risky Choices out-of-context: A Monte Carlo Study." Technical report, University of Houston.

Wilcox, N. (2011), "Stochastically more Risk Averse: A Contextual Theory of Stochastic Discrete Choice under Risk." *Journal of Econometrics*, 162, 89 – 104. The Economics and Econometrics of Risk.

Wright, J. and K. Leyton-Brown (2017), "Predicting Human Behavior in Unrepeated, Simultaneous-move Games." *Games and Economic Behavior*, 16–37.

Zauner, G., K. (1999), "A Payoff Uncertainty Explanation of Results in EExperimental Centipede Games." *Games and Economic Behavior*, 26, 157–185.

Zhang, P. (1993), "Model Selection via Multifold Cross Validation." *The Annals of Statistics*, 162, 299–313.

Appendix A Simulation (for online publication)

To test whether the models are properly identified and are indeed behaviourally distinguishable, we conduct an extensive Monte Carlo simulation with two parts similar to the one in Wilcox (2006) and Hey and Pace (2014). For each of the models, we assume a reasonable set of parameters and simulate experimental datasets. As all of the specifications incorporate a deterministic choice rule within a stochastic model, the parameters of the stochastic part reflect the level of noise in the subjects' decisions. When the level of noise is very low, the subjects behave according to the Nash equilibrium prediction (play T at the first node). As the level of noise increases, the statistical models predict that players will end the games at a later node. To this end, it makes sense to assume different levels of noise ranging from low to high. In particular we assume three levels of noise, a LOW level, where the predicted outcome is that the majority of the subjects ends the game during the first three nodes, a MEDIUM level, where the game ends in nodes 3-5, which is also what is usually observed in experimental studies, and a HIGH level, where the predicted probability is more uniformly distributed between all the nodes of the game.

Table 7 reports the values that were used for each of the models. We also use an additional set of parameters, the ACTUAL, which corresponds to the estimated values obtained by our analysis reported in Table 2. Using these simulated data, we generate predicted probabilities. To generate the datasets, we follow Wilcox (2006) and introduce variability by drawing the parameters from a Gamma distribution, particularly specified for each parameter, which is centered to the assumed values within each set. We set a level of variability equal to 5% (coefficient of variation)[23]. For example, for the EU-AQRE specification, the MEDIUM λ parameter for each simulation is drawn from a Gamma distribution with shape and rate values equal to 20 (these values correspond to mean equal to 1 and standard deviation 0.05). For the param-

[23]The results are qualitatively the same for various levels of variability. 5% level generates enough variability for a meaningful simulation.

eters which are common between models, such as the weighting function coefficient γ or the power utility coefficient r we assumed values that are often observed in the literature. More specifically, we assume that $\gamma = 0.61$, $r = 0.70$, $\tau = 1.50$ for the LEVEL-K model and $\rho = 0.01$ for the AQRE+[24].

The simulation exercise proceeds as follows. For a given model under consideration, we generate a simulated dataset. Then, this dataset is fitted to generate estimates for all the 14 models along with goodness of fit measures. If it is possible to identify the data generating model, this will be demonstrated by a goodness of fit measure (the data generating model should provide a better fit compared to the other models). We test this by comparing the value of the Akaike Information Criterion (AIC) between all models, where the best fitting model should generate the lowest value of AIC. We repeat the same procedure for 100 simulations. Tables 9-12 report the average value of AIC for all cross-model comparisons for the LOW, MEDIUM, HIGH and ACTUAL sets of parameters respectively. Each row model corresponds to the data-generating model (the true model) and each column model corresponds to the one used during the estimation. The results in the Tables show that the models under analysis are fully distinguishable. In each row, the element of the diagonal (when the true and the assumed model coincide) is always smaller than the off-diagonal elements. That is for instance, if we know that the data-generating model is the RDU-AQRE, then the RDU-AQRE model, better fits the data. For example, in Table 9, when both the true and the estimated model is the RDU-AQRE, the average value of AIC is equal to 808.33, while when the true model is the RDU-AQRE and the estimated is the EU-AQRE, the average AIC is 844.11, indicating that when the true and the estimated models are the same, we obtain a better fit and therefore we can identify the true model. This result is identical across all the sets of simulated parameters, LOW, MEDIUM and HIGH, as well as for the ACTUAL set of parameters.

The second part of the simulation was to confirm that it is not only possible to identify

[24] Introducing variability to these parameters does not change the results. We keep these parameters fixed in order to focus on those parameters that differ between models.

the data generating model, but also to recover the value of the parameters that generated the simulated dataset. Table 8 reports the mean estimated value of the parameters for all the specifications based on the 100 simulations, when the true and the assumed model used for the estimation coincide. The first column reports the actual value of the parameters (MEDIUM set), while columns 2-5 report the mean estimated parameters of the models. Below each estimated parameter, the standard deviation is reported. From the Table it is apparent that the estimated value of the parameters is, on average, quite close to the actual value. The relatively large standard deviation confirms that despite the variability in the simulated datasets, it is still possible to identify the data-generating model and recover the value of the parameters.

Model	LOW	MEDIUM	HIGH	ACTUAL
EU-AQRE	λ=0.5	λ=1	λ=2	λ=2.131
RDU-AQRE	λ=0.5	λ=1	λ=2	λ=1.719
RDU-AQRE-U	λ=0.1	λ=0.15	λ=0.2	λ=0.07
LEVEL-K	λ=0.05	λ=0.1	λ=0.2	λ=0.018
AQRE+	λ=0.5	λ=1	λ=2	λ=1.712

Table 7: The Table reports the values of the three set of parameters (LOW, MEDIUM, HIGH and ACTUAL) that were used for the simulation. The parameters differ between models since they represent error in choices in different ways each. For the shared parameters we set $\gamma = 0.61, r = 0.70, \tau = 1.5$ and $\rho = 0.01$.

Model	REAL	ESTIMATED		
EU-AQRE	λ=1	λ=1.003		
		0.057		
RDU-AQRE	λ=1	$\lambda = 1.003$	$\gamma = 0.610$	
		0.096	0.07	
RDU-AQRE-U	λ=0.15	$\lambda = 0.149$	$\gamma = 0.642$	$r = 0.682$
		0.154	0.375	0.286
LEVEL-K	λ=0.1	λ=0.092	$\tau = 1.501$	
		0.397	0.08	
AQRE+	λ=1	λ=0.991	$\rho = 0.011$	
			0.02	

Table 8: The Table reports the actual value of the parameters (MEDIUM set) used for the simulation (column 1) along with the mean value of the estimated parameters (columns 2-5) from 100 simulations for each of the models. The standard deviation is reported below each estimate. The shared parameters of probability weighting γ and risk aversion r for the RDU specifications, have been set to 0.61 and 0.70 respectively. The parameter τ for the LEVEL-K model is set to 1.5 and the ρ parameter of the AQRE+ model has been set to 0.01.

	EU-AQRE	RDU-AQRE	RDU-AQRE-U	LEVEL-K	AQRE+	MIN(AIC)
EU-AQRE	**786.21**	787.86	792.15	958.19	800.26	786.21
RDU-AQRE	844.11	**808.33**	813.88	970.41	820.97	808.33
RDU-AQRE-U	830.06	779.68	**732.69**	1030.86	882.16	732.69
LEVEL-K	1328.83	1219.43	1108.41	**1017.22**	1221.22	1017.22
AQRE+	986.28	787.62	785.12	826.63	**650.95**	650.95

Table 9: The Table reports the results from the LOW simulation. Each cell reports the mean Akaike Information Criterion (the lower the better) from 100 simulations, for each column model when the row model is the data-generating one (true model). The last column reports the minimum average value of the AIC for each cross-model comparison for each specification.

	EU-AQRE	RDU-AQRE	RDU-AQRE-U	LEVEL-K	AQRE+	MIN(AIC)
EU-AQRE	**939.04**	941.04	954.02	1084.78	941.04	939.04
RDU-AQRE	1005	**935.6**	952.77	1066.51	940.51	935.6
RDU-AQRE-U	1029.49	989.51	**962.74**	1057.82	997.75	962.74
LEVEL-K	1329.52	1219.36	1108.23	**1015.69**	1220.88	1015.69
AQRE+	1040.6	896	899.82	969.96	**854.88**	854.88

Table 10: The Table reports the results from the MEDIUM simulation. Each cell reports the mean Akaike Information Criterion (the lower the better) from 100 simulations, for each column model when the row model is the data-generating one (true model).The last column reports the minimum average value of the AIC for each cross-model comparison for each specification.

	EU-AQRE	RDU-AQRE	RDU-AQRE-U	LEVEL-K	AQRE+	MIN(AIC)
EU-AQRE	**1039.43**	1041.43	1062.22	1219.74	1041.43	1039.43
RDU-AQRE	1130.07	**1061.79**	1094.14	1266.02	1063.66	1061.79
RDU-AQRE-U	1060.89	1023.33	**998.24**	1114.91	1048.69	998.24
LEVEL-K	1333.11	1218.92	1107.57	**1007.99**	1219.01	1007.99
AQRE+	1100.64	1056.47	1084.82	1249.28	**1055.55**	1055.55

Table 11: The Table reports the results from the HIGH simulation. Each cell reports the mean Akaike Information Criterion (the lower the better) from 100 simulations, for each column model when the row model is the data-generating one (true model). The last column reports the minimum average value of the AIC for each cross-model comparison for each specification.

	EU-AQRE	RDU-AQRE	RDU-AQRE-U	LEVEL-K	AQRE+	MIN(AIC)
EU-AQRE	**1043.84**	1045.84	1066.82	1230.26	1045.84	1043.84
RDU-AQRE	1103.55	**1028.21**	1058.02	1207.25	1030.62	1028.21
RDU-AQRE-U	586.31	511.09	**425.57**	1018.39	739.75	425.57
LEVEL-K	1328.71	1219.36	1108.96	**1017.97**	1221.09	1017.97
AQRE+	1090.47	1025.9	1049.39	1187.26	**1022.78**	1022.78

Table 12: The Table reports the results from the ACTUAL simulation. Each cell reports the mean Akaike Information Criterion (the lower the better) from 100 simulations, for each column model when the row model is the data-generating one (true model). The last column reports the minimum average value of the AIC for each cross-model comparison for each specification.