

Ambiguity and the Centipede Game: Strategic Uncertainty in Multi-Stage Games with Almost Perfect Information.*

Jürgen Eichberger
Alfred Weber Institut,
Universität Heidelberg, Germany.

Simon Grant
Research School of Economics,
Australian National University, Australia.

David Kelsey
Department of Economics,
University of Exeter, England.

April 23, 2018

Abstract

We propose a solution concept, consistent-planning equilibrium under ambiguity (CP-EUA), for two-player multi-stage games with almost perfect information. Players are neo-expected payoff maximizers. The associated (ambiguous) beliefs are revised by Generalized Bayesian Updating. Individuals take account of possible changes in their preferences by using consistent planning. We show that if there is ambiguity in the centipede game and players are sufficiently optimistic then it is possible to sustain ‘cooperation’ for many periods. Similarly, in a non-cooperative bargaining game we show that there may be delay in agreement being reached.

Keywords: optimism, neo-additive capacity, extensive-form games, dynamic consistency, consistent planning, centipede game.

JEL classification: D81

*For comments and suggestions we would like to thank Pierpaolo Battigalli, Lorenz Hartmann, Philippe Jehiel, David Levine, George Mailath, Larry Samuelson, Marciano Siniscalchi, Rabee Tourky, Qizhi Wang, participants in seminars at ANU, Bristol and Exeter, FUR (Warwick 2016), DTEA (Paris 2017) and SAET (Faro 2017).

1 Introduction

More than 70 years after von Neumann and Morgenstern's (1944) path-breaking "Theory of Games and Economic Behavior" and more than a decade after the Nobel Prize in Economics had been awarded to John Nash, John Harsanyi and Reinhard Selten for their contributions to game theory, laboratory experiments, econometric studies and field work have amply demonstrated that actual human behavior in interactive situations diverges significantly from Nash equilibrium. Behavioral economics has replaced economic theory as the most dynamic branch of economics. Yet, despite its success in recording numerous robust deviations from "rational behavior," it has failed to provide a new paradigm of how to model economic interaction.

From its beginning, game theory had been intimately related to decision making under uncertainty. Indeed, von Neumann and Morgenstern suggested the axiomatic approach to expected utility in order to provide a foundation for the evaluation of mixed strategies. Yet, while decision theory responded to the behavioral challenges of the Allais and Ellsberg paradoxes with a plethora of alternative preference representations, suitable for accommodating these and many more behavioral biases observed in experiments, game theory retains its early paradigm of rational expectations in the sense of players perfectly predicting the probability distributions governing the behavior of other players in terms of mixed strategies or any uncertainties about payoffs. The insistence on rational expectations as a necessary requirement for stable interactions has prevented game theory from providing answers to the most obvious challenges for predictions based on Nash equilibrium.¹

The theory of decision making under uncertainty provides a wide range of representations that can accommodate behavioral biases, both for the case when the actual probabilities of events are known, e.g., Prospect Theory (Kahneman and Tversky 1979), and for situations where no or only partial information about the probabilities of events are available, as in Choquet expected utility (Schmeidler 1989) or multiple prior approaches (Gilboa and Schmeidler 1989, Ghirardato, Maccheroni, and Marinacci 2004). Yet, none of these criteria for decision making under ambiguity has been successfully implemented in a game-theoretic context where uncertainty concerns the behavior of the opponents.² In our opinion, the main obstacles have been the difficult questions regarding:

¹There 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 would allow one to make ex ante predictions. We will discuss this literature in Section 7 in more detail.

²There is a small literature on ambiguity in games in strategic form. Our earlier approach, Eichberger and Kelsey (2000), has its roots in Dow and Werlang (1994) and is similar to Marinacci (2000).

More recently, Riedel and Sass (2014) explore more general strategy notions ("Ellsberg strategies") in games of

- how much consistency should be required between the beliefs of players about their opponents' behavior and their actual behavior in order to justify calling a situation an (at least temporary) equilibrium; and,
- how much consistency to impose on dynamic strategies since all consequentialist updating rules for non-expected-utility representations essentially entail violations of dynamic consistency.³

In this paper, we suggest a general notion of equilibrium (Equilibrium under Ambiguity) studied in Eichberger and Kelsey (2014) and a general notion of updating for beliefs (Generalized Bayesian Updating) analysed in Eichberger, Grant, and Kelsey (2007). Both concepts are general in that they can be applied to preference representations where beliefs are represented by capacities (Choquet expected utility, prospect theory) or by sets of multiple priors (Maxmin and α -maxmin expected utility). Applying these concepts to the *non extreme outcome (neo)*-expected utility representation axiomatized in Chateauneuf, Eichberger, and Grant (2007), allows us, with just two additional (unit-interval valued) parameters, to model a player's behavior in the face of strategic ambiguity. The first parameter reflects the player's perception of strategic ambiguity by measuring that player's degree of confidence regarding their probabilistic beliefs about their opponent's behavior. The second measures the player's relative pessimistic versus optimistic attitudes toward this perceived strategic ambiguity.⁴ Probabilistic beliefs are endogenously determined in equilibrium and are updated in the usual Bayesian way when new information arrives. With new information, however, a player's degree of confidence will change as well and, hence, the impact of a player's relative pessimistic versus optimistic attitudes toward the strategic ambiguity. This novel framework allows us to study the role of strategic ambiguity in games both under optimism and pessimism.⁵

In order to ensure dynamic consistency, we will adapt the notion of consistent planning proposed by Strotz (1955) and Siniscalchi (2011) for the game-theoretic context.⁶ As we argue below, consistent

complete information and Azrieli and Teper (2011), Kajii and Ui (2005) and Grant, Meneghel, and Tourky (2016) study games of incomplete information with ambiguity about priors.

³An exception is the recursive multiple priors model of Epstein and Schneider (2003) that retains dynamic consistency with a consequentialist updating rule but at the cost of imposing stringent restrictions on what form the information structure may take.

⁴Neo-expected utility is a special case of Choquet expected utility (Schmeidler 1989), of α -multiple priors expected utility (Gilboa and Schmeidler 1989, Ghirardato, Maccheroni, and Marinacci 2004), and of rank-dependent expected utility (Quiggin 1982).

⁵There is also a small earlier literature on extensive form games. Lo (1999) provides the first model treating ambiguity in extensive form games. Rothe (2011) proposes a generalization of subgame perfection for players with non-additive beliefs that is similar to the equilibrium concept we develop but for the most part the players in his model only exhibit pessimism toward the ambiguity they perceive there to be about the strategy choice of their opponents. All other papers deal with ambiguity in special cases: Eichberger and Kelsey (1999) and Eichberger and Kelsey (2004) study signaling games and, more recently, Kellner and LeQuement (2015) cheap-talk games and Bose and Renou (2014) mechanism design questions with communication.

⁶? employ a similar notion of consistent planning for their equilibrium with players whose preferences have a multiple

planning finds behavioral support in a large literature in psychology on “self-regulation” (Baumeister and Vohs 2004). Moreover, it allows us to work with the well-known backward induction methodology.

To demonstrate the potential of this new approach, we apply it to multi-stage two-player games with almost perfect information.⁷ It is within this context that many, if not most, deviations of human behavior from Nash equilibrium have been noted. In particular, we show that a small degree of optimism in combination with some ambiguity induces equilibrium behavior which corresponds to behavior observed in experimental studies. As examples, we have chosen two of the most challenging cases from this class of games: the centipede game and the alternating-offer bargaining game. For the former, we provide a complete characterization of equilibria under ambiguity in terms of the perception of ambiguity and attitudes toward perceived ambiguity parameters. For the latter, we show that inefficient delays in bargaining may be the result of ambiguity about the other player’s behavior. Though our framework allows us to derive equilibria by adapting well-known backward induction methods, the results reveal new channels of influence on behavior. In particular, the importance of some optimism in the face of uncertainty is highlighted, a channel of influence widely disregarded, since almost all preference representations under ambiguity have been axiomatized and analysed for the case of pessimism only.

1.1 Backward Induction and Ambiguity

The standard analysis of sequential two-player games with complete and perfect information uses backward induction or subgame perfection in order to rule out equilibria which are based on “incredible” threats or promises. In sequential two-player games with perfect information, this principle successfully narrows down the set of equilibria and leads to precise predictions. Sequential bargaining, Rubinstein (1982), repeated prisoner’s dilemma, chain store paradox, Selten (1978), and the centipede game, Rosenthal (1981), provide well-known examples.

Experimental evidence, however, suggests that in all these cases the unique backward induction equilibrium is a poor predictor of behavior.⁸ It appears as if payoffs received off the “narrow” equilibrium path do influence behavior, even if a step by step analysis shows that it is not optimal to deviate from it at any stage, see Greiner (2016). This suggests that we should reconsider the logic of

⁷Osborne and Rubinstein (1994) (p102) refer to this class as extensive games with perfect information and simultaneous moves.

⁸For the bargaining game Güth, Schmittberger, and Schwarze (1982) provided an early experimental study and for the centipede game McKelvey and Palfrey (1992) find evidence of deviations from Nash predictions.

backward induction.

Our concept of *Consistent Planning Equilibrium (in Beliefs) Under Ambiguity* (henceforth CP-EUA) extends the notion of strategic ambiguity to sequential two-player games with perfect information. Despite ambiguity, players remain (sequentially) consistent with regard to their own strategies. With this notion of equilibrium, we reconsider some of the well-known games mentioned above in order to see whether ambiguity about the opponent's strategy brings game-theoretic predictions closer to observed behavior. CP-EUA suggests a general principle for analyzing extensive form games without having to embed them into elaborately structured games of incomplete information.

The notion of an Equilibrium under Ambiguity (EUA) for strategic games in Eichberger and Kelsey (2014) rests on the assumption that players take their knowledge about the opponents' incentives reflected in their payoffs seriously but not as beyond doubt. Although they predict their opponents' behavior based on their knowledge about the opponents' incentives, they do not have full confidence in these predictions. There may be very little ambiguity if the interaction takes place in a known context with familiar players or it may be substantial in unfamiliar situations where the opponents are strangers. In contrast to standard Nash equilibrium theory, in an EUA the cardinal payoffs of a player's own strategies may matter if they are particularly high (optimistic case) or particularly low (pessimistic case). Hence, there will be a trade off between relying on the prediction about the opponents' behavior and the salience of one's own strategy in terms of the outcome.

In dynamic games, where a strategy involves a sequence of moves, the observed history may induce a reconsideration of previously planned actions. As a result, the analysis needs to consider issues of dynamic consistency and also whether equilibria rely upon incredible threats or promises. The logic of backward induction forbids a player to consider any move of the opponent which is not optimal, no matter how severe the consequences of such a deviation may be. This argument is weaker in the presence of ambiguity. In contrast in our equilibrium, players maintain sophistication by having correct beliefs about their own future moves.

These considerations suggest that ambiguity makes it harder to resolve dynamic consistency problems. However there are also advantages to studying ambiguous beliefs. We shall update beliefs by the commonly used Generalized Bayesian Updating rule (henceforth GBU). It has the property that it is usually defined both on and off the equilibrium path. This contrasts with standard solution concepts, such as Nash equilibrium or subgame perfection, where beliefs off the equilibrium path are somewhat arbitrary, since Bayes' rule is not defined at such events.

We do not assume that players are solely ambiguity-averse but also allow for optimistic as well as pessimistic attitudes toward the ambiguity which the players perceive. This would imply that the player over-weights both high and low payoffs compared to a standard expected payoff maximizer. As a result middle ranking outcomes are under-weighted. We show that a game of complete and perfect information need not have a pure strategy equilibrium. Thus a well-known property of Nash equilibrium need not apply when there is ambiguity.

1.2 Ambiguity in the Centipede Game

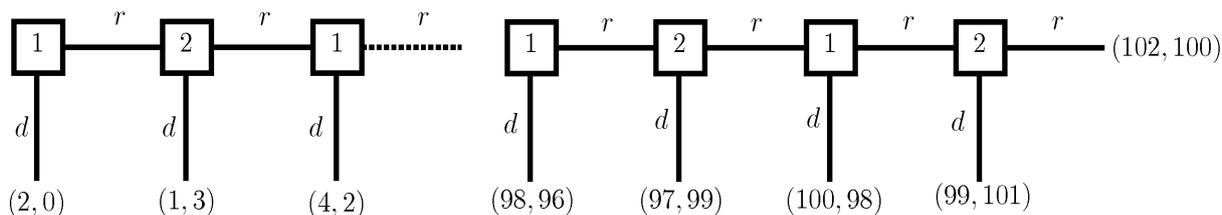


Figure 1: The Centipede Game

The centipede game is illustrated in figure 1. It has been a long-standing puzzle in game theory. Intuition suggests that there are substantial opportunities for the players to cooperate. However standard solution concepts imply that cooperation is not possible. In this game there are two players who move alternately. At each move a player has the option of giving a benefit to her opponent at a small cost to herself. Alternatively she can stop the game at no cost.

Conventional game theory makes a clear prediction. Nash equilibrium and iterated dominance both imply that all equilibria in the centipede game are ones in which the first player to move stops the game by playing down, d . This is despite the fact that both players could potentially make large gains if the game continues until close to the end. Intuition suggests that it is more likely that players will cooperate, at least for a while, thereby increasing both payoffs. This is confirmed by the experimental evidence, see McKelvey and Palfrey (1992).

It is plausible that playing right, r , may be due to optimistic attitudes toward the ambiguity the player perceives there to be about her opponent's choice of strategy. By playing r , a player is choosing between a high but uncertain payoff in preference to a low but safe payoff. One reason why ambiguity may be present in the centipede game, is that many rounds of deletion of dominated strategies are needed to produce the standard prediction. A player may be uncertain as to whether her opponent performs some or all of them.

Our conclusions are that with ambiguity-averse preferences the only equilibrium that remains is the one without cooperation since ambiguity aversion increases the attraction of playing down and receiving a certain payoff. However if players have optimistic attitudes toward ambiguity they may be tempted to cooperate by the high payoffs towards the end of the game. We find that even moderate degrees of ambiguity loving are sufficient to produce cooperation in the centipede game. This is compatible with experimental data on ambiguity-attitudes, for a survey see Trautmann and de Kuilen (2015).

1.3 Bargaining

As a second application we consider non-cooperative bargaining. Sub-game perfection suggests that agreement will be instantaneous and outcomes will be efficient. However these predictions do not seem to be supported in many of the situations which bargaining theory is intended to represent. Negotiation between unions and employers often take substantial periods of time and involve wasteful actions such as strikes. Similarly international negotiations can be lengthy and may yield somewhat imperfect outcomes. We suggest that optimistic attitudes toward ambiguity might play a role in explaining this. Parties to a bargain initially choose ambitious positions in the hope of achieving large gains. If these expectations are not realized they later shift to make more reasonable demands.

Organization of the paper We first describe in section 2, how we model ambiguity and the rule we use for updating as well as our approach to dynamic choice. In section 3 we present the class of games we shall be studying along with the attendant notation. We then explain how we incorporate into these games the model of ambiguity developed in the previous section. In section 4 we present our solution concept. We demonstrate existence and show that games of complete and perfect information may not have pure equilibria. This is applied to the centipede game in section 5 and to bargaining in section 6. The related literature is discussed in section 7 and section 8 concludes. The appendix contains proofs of those results not proved in the text.

2 Framework and Definitions

In this section we describe how we model ambiguity, updating and dynamic choice.

2.1 Ambiguous Beliefs and Expectations

For a typical two-player game let $i \in \{1, 2\}$, denote a generic player. We shall adopt the convention of referring to player 1 (respectively, player 2) by female (respectively, male) pronouns and a generic player by plural pronouns. Let S_i and S_{-i} denote respectively the finite strategy sets of player i and that of their opponent. We denote the payoff to player i from choosing their strategy s_i in S_i , when their opponent has chosen s_{-i} in S_{-i} by $u_i(s_i, s_{-i})$. Following Schmeidler (1989) we shall model ambiguous beliefs of player i on S_{-i} with a particular sub-class of capacities, where a capacity is a monotonic and normalized set function.

Definition 2.1 *A capacity on S_{-i} is a real-valued function ν_i on the subsets of S_{-i} such that $A \subseteq B \Rightarrow \nu_i(A) \leq \nu_i(B)$ and $\nu_i(\emptyset) = 0$, $\nu_i(S_{-i}) = 1$. Moreover, the capacity is convex (respectively, concave) if for all $A, B \subseteq S_{-i}$, $\nu_i(A) + \nu_i(B) \leq$ (respectively, \geq) $\nu_i(A \cap B) + \nu_i(A \cup B)$.⁹*

The ‘expected’ payoff associated with a given strategy s_i in S_i of player i , with respect to the capacity ν_i on S_{-i} is taken to be the Choquet integral, defined as follows.

Definition 2.2 *The Choquet integral of $u_i(s_i, \cdot)$ with respect to the capacity ν_i on S_{-i} is:*

$$V_i(s_i | \nu_i) = \int u_i(s_i, s_{-i}) d\nu_i(s_{-i}) = u_i(s_i, s_{-i}^1) \nu(s_{-i}^1) + \sum_{r=2}^R u_i(s_i, s_{-i}^r) [\nu_i(s_{-i}^1, \dots, s_{-i}^r) - \nu_i(s_{-i}^1, \dots, s_{-i}^{r-1})],$$

where $R = |S_{-i}|$ and the strategy profiles in S_{-i} are numbered so that $u_i(s_i, s_{-i}^1) \geq u_i(s_i, s_{-i}^2) \geq \dots \geq u_i(s_i, s_{-i}^R)$.

Preferences represented by a Choquet integral with respect to a capacity are referred to as Choquet Expected Utility (henceforth CEU).

2.2 Neo-additive capacities

As argued in the introduction, capacities as defined in Definition 2.1 and their CEU are far too general for applications in economics and game theory. Hence, we will restrict attention to a special class of capacities with a small number of parameters which have natural interpretations, a simple CEU, and intuitive notions of updating. Chateauneuf, Eichberger, and Grant (2007) have axiomatized a

⁹A probability measure is the special case of a capacity that is both convex and concave, that is, it is additive: $\nu_i(A) + \nu_i(B) = \nu_i(A \cap B) + \nu_i(A \cup B)$.

parsimoniously parametrized special case of CEU, that we shall refer to as *non extreme outcome (neo)-expected payoff* preferences. Capacities in this sub-class of CEU are characterized by a probability distribution π_i on S_{-i} and two additional parameters $\alpha_i, \delta_i \in [0, 1]$.

Definition 2.3 *The neo-additive capacity is defined by setting:*

$$\nu_i(A|\alpha_i, \delta_i, \pi_i) = \begin{cases} 1 & \text{for } A = S_{-i} \\ \delta_i(1 - \alpha_i) + (1 - \delta_i)\pi_i(A) & \text{for } \emptyset \neq A \subset S_{-i} \\ 0 & \text{for } A = \emptyset \end{cases} .$$

Fixing the two parameters α_i and δ_i , it is straightforward to show for any probability distribution π_i on S_{-i} , that the Choquet integral of $u_i(s_i, \cdot)$ with respect to the neo-additive capacity $\nu_i(\cdot|\alpha_i, \delta_i, \pi_i)$ on S_{-i} takes the simple and intuitive form of a weighted average between the expected utility with respect to π_i and the α -maxmin utility suggested by Hurwicz (1951) for choice under complete ignorance.

Lemma 2.1 *The CEU with respect to the neo-additive capacity $\nu_i(\cdot|\alpha_i, \delta_i, \pi_i)$ on S_{-i} can be expressed as:*

$$V_i(s_i|\nu_i(\cdot|\alpha_i, \delta_i, \pi_i)) = (1 - \delta_i) \cdot \mathbf{E}_{\pi_i} u_i(s_i, \cdot) + \delta_i \left[\alpha_i \min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) + (1 - \alpha_i) \max_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) \right], \quad (1)$$

where $\mathbf{E}_{\pi_i} u_i(s_i, \cdot)$ denotes a conventional expectation taken with respect to the probability distribution π_i on S_{-i} .¹⁰

One can interpret π_i as player i 's "probabilistic belief" or "theory" about how their opponent is playing. However, players perceive there to be some degree of ambiguity associated with their theory. Their "confidence" in this theory is modelled by the weight $(1 - \delta_i)$ given to the expected payoff $\mathbf{E}_{\pi_i} u_i(s_i, \cdot)$. Correspondingly, the highest (respectively, lowest) possible degree of ambiguity corresponds to $\delta_i = 1$, (respectively, $\delta_i = 0$). Their attitude toward such ambiguity is measured by α_i . Purely ambiguity-loving preferences are given by $\alpha_i = 0$, while the highest level of ambiguity-aversion occurs when $\alpha_i = 1$. If $0 < \alpha_i < 1$, the player has a mixed attitude toward ambiguity, since they respond to ambiguity partly in an pessimistic way by over-weighting the worst feasible payoff and partly in an optimistic way by over-weighting the best feasible payoff associated with that strategy.

¹⁰To simplify notation we shall suppress the arguments and write $V_i(s_i|\nu_i)$ for $V_i(s_i|\nu_i(\cdot|\alpha_i, \delta_i, \pi_i))$ when the meaning is clear from the context.

2.3 The Support of a Capacity

We wish the support of a capacity to represent those strategies that a given player believes their opponent might play. Sarin and Wakker (1998) argue that the decision-maker's beliefs may be deduced from the decision weights in the Choquet integral. With this in mind, we propose the following definition.

Definition 2.4 Fix ν_i a capacity on S_{-i} . Set

$$\text{supp}(\nu_i) := \{s_{-i} \in S_{-i} : \forall A \subset S_{-i} \setminus \{s_{-i}\}; \nu_i(A \cup s_{-i}) > \nu_i(A)\}.$$

Notice that the set $\text{supp}(\nu_i)$ consists of those strategies of i 's opponent which always receive positive weight in the Choquet integral, no matter which of i 's strategies is being evaluated. To see this, recall that the Choquet expected payoff of a given strategy s_i , is a weighted sum of payoffs for which the decision-weight assigned to their opponent's strategy \tilde{s}_{-i} is given by

$$\nu_i(\{s_{-i} : u_i(s_i, s_{-i}) > u_i(s_i, \tilde{s}_{-i})\} \cup \{\tilde{s}_{-i}\}) - \nu(\{s_{-i} : u_i(s_i, s_{-i}) > u_i(s_i, \tilde{s}_{-i})\}).$$

These weights depend on the way in which the strategy s_i ranks the strategies in S_{-i} . Since there are $R!$ ways the elements of S_{-i} can be ranked, in general there are $R!$ decision weights used in evaluating the Choquet integral with respect to a given capacity.

The next lemma shows that $\text{supp}(\nu_i)$ yields an intuitive result when applied to neo-additive capacities.

Lemma 2.2 Let ν_i be a neo-additive capacity on S_{-i} , then $\text{supp}(\nu_i) = \text{supp} \pi_i = \{s_{-i} \in S_{-i} : \pi_i(s_{-i}) > 0\}$.

Proof. If $s_{-i} \notin A \subset S_{-i}$, then $\nu_i(A \cup \{s_{-i}\}) - \nu_i(A) = [\delta_i(1 - \alpha_i) + (1 - \delta_i)\pi_i(A \cup s_{-i})] - [\delta_i(1 - \alpha_i) + (1 - \delta_i)\pi_i(A)] = [(1 - \delta_i)\pi_i(A) + (1 - \delta_i)\pi_i(s_{-i})] - [(1 - \delta_i)\pi_i(A)] = (1 - \delta_i)\pi_i(s_{-i})$.

Thus $\nu_i(A \cup \{s_{-i}\}) > \nu_i(A) \Leftrightarrow \pi_i(s_{-i}) > 0$. ■

Recall that we interpret π_i as the ‘‘probabilistic belief’’ or ‘‘theory’’ of an ambiguous belief. Thus it is natural that the support of the capacity ν_i is the support in the usual sense of the additive probability π_i .

2.4 Updating Ambiguous Beliefs

CEU is a theory of decision-making at one point in time. To use it in extensive form games we need to extend it to multiple time periods. We do this by employing Generalized Bayesian Updating (henceforth GBU) to revise beliefs. One problem which we face is that the resulting preferences may not be dynamically consistent. We respond to this by assuming that individuals take account of future preferences by using *consistent planning*, defined below. The GBU rule has been axiomatized in Eichberger, Grant, and Kelsey (2007) and Horie (2013). It is defined as follows.

Definition 2.5 *Let ν_i be a capacity on S_{-i} and let $E \subseteq S_{-i}$. The Generalized Bayesian Update (henceforth GBU) of ν_i conditional on E is given by:*

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

where $E^c = S_{-i} \setminus E$ denotes the complement of E .

The GBU rule coincides with Bayesian updating when beliefs are additive.

Lemma 2.3 *For a neo-additive belief $\nu_i(\cdot | \alpha_i, \delta_i, \pi_i)$ the GBU conditional on E is given by*

$$\nu_i^E(A | \alpha_i, \delta_i, \pi_i) = \begin{cases} 0 & \text{if } A \cap E = \emptyset, \\ \delta_i^E (1 - \alpha_i) + (1 - \delta_i^E) \pi_i^E(A) & \text{if } \emptyset \subset A \cap E \subset E, \\ 1 & \text{if } A \cap E = E. \end{cases}$$

where $\delta_i^E = \delta_i / [\delta_i + (1 - \delta_i)\pi_i(E)]$, and $\pi_i^E(A) = \pi_i(A \cap E) / \pi_i(E)$.

Notice that for a neo-additive belief with $\delta_i > 0$, the GBU update is well-defined even if $\pi_i(E) = 0$ (that is, E is a zero-probability event according to the individual's 'theory'). In this case the updated parameter $\delta_i^E = 1$, which implies the updated capacity is a Hurwicz capacity that assigns the weight 1 to every event that is a superset of E , and $(1 - \alpha)$ to every event that is a non-empty strict subset of E .

The following result states that a capacity is neo-additive, if and only if both it and its GBU update admit a multiple priors representation with the same α_i and the updated set of beliefs is the prior by prior Bayesian update of the initial set of probabilities.¹¹

¹¹A proof can be found in Eichberger, Grant, and Kelsey (2012).

Proposition 2.1 *The capacity ν_i on S_{-i} is neo-additive for some parameters α_i and δ_i and some probability π_i , if and only if both the ex-ante and the updated preferences respectively admit multiple priors representations of the form:*

$$\int u_i(s_i, s_{-i}) d\nu_i(s_{-i}) = \alpha_i \times \min_{q \in \mathcal{P}} \mathbf{E}_q u_i(s_i, \cdot) + (1 - \alpha_i) \times \max_{q \in \mathcal{P}} \mathbf{E}_q u_i(s_i, \cdot),$$

$$\int u_i(s_i, s_{-i}) d\nu_i^E(s_{-i}) = \alpha_i \times \min_{q \in \mathcal{P}^E} \mathbf{E}_q u_i(s_i, \cdot) + (1 - \alpha_i) \times \max_{q \in \mathcal{P}^E} \mathbf{E}_q u_i(s_i, \cdot),$$

where $\mathcal{P} := \{p \in \Delta(S_{-i}) : p \geq (1 - \delta_i) \pi_i\}$, $\mathcal{P}^E := \{p \in \Delta(E) : p \geq (1 - \delta_i^E) \pi_i^E\}$, $\delta_i^E = \frac{\delta_i}{\delta_i + (1 - \delta_i) \pi_i(E)}$, and $\pi_i^E(A) = \frac{\pi_i(A \cap E)}{\pi_i(E)}$.

We view this as a particularly attractive and intuitive result since the ambiguity-attitude, α_i , can be interpreted as a characteristic of the individual which is *not updated*. In contrast, the set of priors is related to the environment and one would expect it to be revised on the receipt of new information.¹²

2.5 Consistent Planning

As we have already foreshadowed, the combination of CEU preferences and GBU updating is not, in general, dynamically consistent. Perceived ambiguity is usually greater after updating. Thus for an ambiguity-averse individual, constant acts will become more attractive. Hence if an individual is ambiguity-averse, in the future she may wish to take a option which gives a certain payoff, even if it was not in her original plan to do so. Following Strotz (1955), Siniscalchi (2011) argues against *commitment to a strategy* in a sequential decision problem in favour of *consistent planning*. This means that a player takes into account any changes in their own preference arising from updating at future nodes. As a result, players will take a sequence of moves which is consistent with backward induction. In general it will differ from the choice a player would make at the first move with commitment.¹³ With consistent planning, however, dynamic consistency is no longer an issue.¹⁴ The dynamic consistency issues and consistent planning are illustrated by the following example of individual choice in the presence of sequential resolution of uncertainty.

Example 2.1 *Consider the following setting of sequential resolution of uncertainty. There are three*

¹²There are two alternative rules for updating ambiguous beliefs, the Dempster-Shafer (pessimistic) updating rule and the Optimistic updating rule, Gilboa and Schmeidler (1993). However neither of these will leave ambiguity-attitude, α_i , unchanged after updating. The updated α_i is always 1 (respectively, 0) for the Dempster-Shafer (respectively, Optimistic) updating rule. See Eichberger, Grant, and Kelsey (2010). For this reason we prefer the GBU rule.

¹³From this perspective, commitment devices should be explicitly modeled. If a commitment device exists, e.g., handing over the execution of a plan to a referee or writing an enforceable contract, then no future choice will be required.

¹⁴Bose and Renou (2014) and Karni and Safra (1989) use versions of consistent planning in games.

time periods, $t = 0, 1, 2$. In period 0 the decision-maker decides whether or not to accept a bet b_W which pays 1 in the event W (Win) and 0 in the complementary event L (Lose). The alternative is to choose an act \bar{b} which yields a certain payoff of x , $0 < x < 1$. At time $t = 1$ she receives a signal which is either good G or bad B . A good (respectively, bad) signal increases (respectively, decreases) the likelihood of winning. If at time 0 she chose to bet and the signal is good she now has the option of switching to a certain payment b_G . In effect selling her bet. To summarize at time $t = 0$ the individual can choose among the following three ‘strategies’:

\bar{b} accept a non-state contingent (that is, guaranteed) payoff of x ; or,

b_G accept the bet but switch to a certain payment if the signal at $t = 1$ is good or

b_W accept the bet and retain it in period 1.

Suppose that the individual is a neo-expected payoff maximizer with capacity ν . Her ‘probabilistic belief’ about the data generating process can be summarized by the following three probabilities: $\pi(G) = p$, $\pi(W|G) = q$ and $\pi(W|B) = 0$, where $\max\{p, q\} < 1$ and $\min\{p, q\} > 0$. Her ‘lack of confidence’ in her belief is given by the parameter $\delta \in (0, 1)$, and her attitude toward ambiguity is given by the parameter α which we assume lies in the interval $(1 - q, 1)$.

The strategy b_G yields a constant payoff of q if the signal realization is G , while b_W leads to a payoff of 1 if the event W obtains and 0 otherwise. The state-contingent payoffs associated with these three strategies are given in the following matrix.

		Events		
		B	$G \cap L$	$G \cap W$
Bets	\bar{b}	x	x	x
	b_G	0	q	q
	b_W	0	0	1

One-shot resolution If the individual is not allowed to revise her choice after learning the realization of the signal in period 1 (or she can commit not to revise her choice), then the choice between b_G and b_W is governed by her ex ante preferences which we take to be represented by the neo-expected payoffs:

$$V(b_G|\nu(\cdot|\alpha, \delta, \pi)) = (1 - \delta)pq + \delta(1 - \alpha)q \text{ and } V(b_W|\nu(\cdot|\alpha, \delta, \pi)) = (1 - \delta)pq + \delta(1 - \alpha).$$

Notice that $V(b_W|\nu) - V(b_G|\nu) = \delta(1 - \alpha)(1 - q) > 0$. Furthermore, if x is set so that $x = (1 - \delta)pq + \delta(1 - \alpha)(1 + q)/2$, then we also have

$$V(\bar{b}|\nu(\cdot|\alpha, \delta, \pi)) = \frac{1}{2}V(b_G|\nu(\cdot|\alpha, \delta, \pi)) + \frac{1}{2}V(b_W|\nu(\cdot|\alpha, \delta, \pi)).$$

So for a one-shot resolution scenario, the individual will strictly prefer to choose b_W over both \bar{b} and b_G .

Sequential resolution Now consider the scenario in which the individual has the opportunity to revise her choice after she has learned the realization of the signal. Let ν^G (respectively, ν^B) denote the GBU of ν conditional on G (respectively, B) realizing. If the signal realization is B then she is indifferent between the pair of bets b_W and b_G . However consider the case where she learns the signal realization is G . According to her *ex ante preferences*, she should stay with her choice of b_W . However, her *updated preference* between the two bets b_W and b_G are governed by the pair of neo-expected payoffs

$$V^G(b_G|\nu^G) = q \quad \text{and} \quad V^G(b_W|\nu^G) = (1 - \delta^G)q + \delta^G(1 - \alpha), \quad \text{where} \quad \delta^G = \frac{\delta}{\delta + (1 - \delta)p}.$$

Notice that $V^G(b_G|\nu^G) - V^G(b_W|\nu^G) = \frac{\delta}{\delta + (1 - \delta)p} \times (\alpha - [1 - q]) > 0$. Hence we have

$$V^G(b_G|\nu^G) > V^G(b_W|\nu^G) \quad \text{and} \quad V^B(b_G|\nu^B) = V^B(b_W|\nu^B) (= 0),$$

$$\text{but } V(b_W|\nu) > V(b_G|\nu),$$

a violation of dynamic consistency (or what Skiadas (1997) calls “coherence”).

Naive Choice versus Consistent Planning If the individual is “naive” then in the sequential resolution setting, she does not choose \bar{b} in period 1, planning to go with b_W in the event the signal realization is G . However, given her updated preferences, she changes her plan of action and chooses the bet b_G instead, yielding her a now guaranteed payoff of q . On the other hand, a consistent planner, anticipating her future self would choose not to remain with the bet b_W after learning the realization of the signal was G , understands that her choice in the first period is really between \bar{b} and b_G . Hence she selects \bar{b} , since $V(\bar{b}|\nu) > V(b_G|\nu)$.

Remark 2.1 *Consistent Planning has also a behavioral component which psychologists relate to as the volitional control of emotions. Optimism and pessimism may be viewed as emotional responses to uncertainty which cannot be “quantified” by the frequency of observations. Being aware of such biases may stimulate “self control” and “self-regulation”. Eisenberg, Smith, and Spinrad (2004, p. 263) write: “Effortful control pertains to the ability to wilfully or voluntarily inhibit, activate, or change (modulate) attention and behavior, as well as executive functioning tasks of planning, detecting errors, and integrating information relevant to selecting behavior.” Taking control of one’s predictable biases, as it is suggested by consistent planning, is probably an essential property of human decision makings. Hence, consistent planning seems to be the adequate self-regulation strategy against dynamic inconsistencies in an uncertain environment for decision makers aware of their optimistic or pessimistic biases.*

3 Multi-stage Games of Almost Perfect Information

We turn now to a formal description of the sequential strategic interaction between two decision-makers. This is done by way of multi-stage games that have a fixed finite number of time periods. In any given period the history of previous moves is known to both players. Within a time period simultaneous moves are allowed. We believe these games are sufficiently general to cover many important economic applications entailing strategic interactions.

3.1 Description of the game.

There are 2 players, $i = 1, 2$ and T stages. At each stage t , $1 \leq t \leq T$, each player i simultaneously selects an action a_i^t .¹⁵ Let $a^t = \langle a_1^t, a_2^t \rangle$ denote a profile of action choices by the players in stage t .

The game has a set \overline{H} of histories h which:

1. contains the empty sequence $h^0 = \langle \emptyset \rangle$, (no records);
2. for any non-empty sequence $h = \langle a^1, \dots, a^t \rangle \in \overline{H}$, all subsequences $\hat{h} = \langle a^1, \dots, a^{\hat{t}} \rangle$ with $\hat{t} < t$ are also contained in \overline{H} .

The set of all histories at stage t are those sequences in \overline{H} of length $t - 1$, with the empty sequence h^0 being the only possible history at stage 1. Let H^{t-1} denote the set of possible histories at stage

¹⁵It is without loss of generality to assume that each player moves in every time period. Games where one player does not move at a particular time, say \hat{t} , can be represented by assigning that player a singleton action set at time \hat{t} .

t with generic element $h^{t-1} = \langle a^1, \dots, a^{t-1} \rangle$.¹⁶ Any history $\langle a^1, \dots, a^T \rangle \in \bar{H}$ of length T is a *terminal history* which we shall denote by z . We shall write $Z (= H^T)$ for the subset of \bar{H} that are terminal histories. Let $H = \bigcup_{t=1}^T H^{t-1}$ denote the set of all non-terminal histories and let $\theta = |H|$ denote the number of non-terminal histories.¹⁷ At stage t , all players know the history of moves from stages $\tau = 1$ to $t - 1$.

For each $h \in H$ the set $A^h = \{a \mid (h, a) \in \bar{H}\}$ is called the action set at h . We assume that A^h is a Cartesian product $A^h = A_1^h \times A_2^h$, where A_i^h denotes the set of actions available to player i after history h . The action set, A_i^h , may depend both on the history and the player. A pure strategy specifies a player's move after every possible history.

Definition 3.1 A (pure) strategy of a player $i = 1, 2$ is a function s_i which assigns to each history $h \in H$ an action $a_i \in A_i^h$.

Let S_i denote the strategy set of player i , $S = S_1 \times S_2$, the set of strategy profiles and $S_{-i} = S_j, j \neq i$, the set of strategies of i 's opponent. Following the usual convention, we will sometimes express the strategy profile $s \in S$ as (s_i, s_{-i}) , in order to emphasize that player i is choosing their strategy $s_i \in S_i$ given their opponent is choosing according to the strategy $s_{-i} \in S_{-i}$.

Each strategy profile $s = (s_1, s_2) \in S$ induces a sequence of histories $\langle h_s^1, \dots, h_s^T \rangle$, given by $h_s^1 = (s_1(h^0), s_2(h^0))$ and $h_s^t = \langle h_s^{t-1}, (s_1(h_s^{t-1}), s_2(h_s^{t-1})) \rangle$, for $t = 2, \dots, T$. This gives rise to a collection of functions $\langle \zeta^t \rangle_{t=0}^T$, where $\zeta^0(s) := h^0$ for every strategy profile $s \in S$ and for each $t = 1, \dots, T$, the function $\zeta^t : S \rightarrow H^t$ is recursively constructed by setting $\zeta^t(s) := \langle (s_1(\zeta^{\tau-1}(s)), s_2(\zeta^{\tau-1}(s))) \rangle_{\tau=1}^t$. Each ζ^t is surjective since every history in H^t must arise from some combination of strategies.

A payoff function u_i for player i , assigns a real number to each terminal history $z \in Z$. With a slight abuse of notation, we shall write $u_i(s)$ for the convolution $u_i \circ \zeta^T(s)$. We now have all the elements to define a multi-stage game.

Definition 3.2 A multi-stage game Γ is a triple $\langle \{1, 2\}, \bar{H}, \{u_1, u_2\} \rangle$, where \bar{H} is the set of all histories, and for $i = 1, 2$, u_i characterizes player i 's payoffs.

3.2 Sub-histories, Continuation Strategies and Conditional Payoffs.

A (sub-) history after a non-terminal history $h \in H$ is a sequence of actions h' such that $(h, h') \in \bar{H}$. Adopting the convention that (h, h^0) is identified with h , denote by \bar{H}^h the set of histories following

¹⁶Notice by definition, that $H^0 = \{h^0\}$.

¹⁷Notice that by construction $\bar{H} = H \cup Z$.

h . Let Z^h denote the set of terminal histories following h . That is, $Z^h = \{z' \in \overline{H}^h : (h, z') \in Z\}$.

Consider a given individual, player i , (she). Denote by s_i^h a (continuation-) strategy of player i which assigns to each history $h' \in \overline{H}^h \setminus Z^h$ an action $a_i \in A_i^{(h, h')}$. We will denote by S_i^h the set of all those (continuation-) strategies available to player i following the history $h \in H$ and define $S^h = S_1^h \times S_2^h$ to be the set of (continuation-) strategy profiles. Each strategy profile $s^h = \langle s_1^h, s_2^h \rangle \in S^h$ defines a terminal history in Z^h . Furthermore, we can take $u_i^h : Z^h \rightarrow \mathbb{R}$, to be the payoff function for player i given by $u_i^h(h') = u_i(h, h')$, and correspondingly set $u_i^h(s^h) := u_i^h(h')$ if the continuation strategy profile s^h leads to the play of the sub-history h' . Consider player i 's choice of continuation strategy s_i^h in S_i^h that starts in stage t . To be able to compute her conditional (Choquet) expected payoff, she must use Bayes' Rule to update her theory π_i (a probability measure defined on S_{-i}) to a probability measure defined on S_{-i}^h . In addition it is necessary to update her perception of ambiguity represented by the parameter δ_i . Now, since ζ^{t-1} is a surjection, there exists a well-defined pre-image $S(h) := (\zeta^{t-1})^{-1}(h) \subseteq S$ for any history $h \in H^{t-1}$. The event $S_{-i}(h)$ is the marginal of this event on S_{-i} given by

$$S_{-i}(h) := \left\{ s_{-i} \in S_{-i} : \exists s_i \in S_i, (\zeta^{t-1})^{-1}(s_i, s_{-i}) = h \right\}.$$

Similarly, the event $S_i(h)$ is the marginal of this event on S_i given by

$$S_i(h) := \left\{ s_i \in S_i : \exists s_{-i} \in S_{-i}, (\zeta^{t-1})^{-1}(s_i, s_{-i}) = h \right\}.$$

Suppose that player i 's initial belief about how the opponent is choosing a strategy is given by a capacity ν_i . Then, her evaluation of the Choquet expected payoff associated with her continuation strategy s_i^h is given by:

$$V_i^h(s_i^h | \nu_i) = \int u_i(s_i^h, s_{-i}^h) d\nu_i^h(s_{-i}^h),$$

where ν_i^h is the GBU of ν_i conditional on history h being reached. Hence, in particular, if she is a neo-expected payoff maximizer with $\nu_i = \nu(\alpha_i, \delta_i, \pi_i)$ then her evaluation of the *conditional neo-expected payoff* of her continuation strategy s_i^h is given by:

$$V_i^h(s_i^h | \nu_i^h) = (1 - \delta_i^h) \mathbb{E}_{\pi_i^h} u_i(s_i^h, \cdot) + \delta_i^h \left[\alpha_i \min_{s_{-i}^h \in S_{-i}^h} u_i^h(s_i^h, s_{-i}^h) + (1 - \alpha_i) \max_{s_{-i}^h \in S_{-i}^h} u_i^h(s_i^h, s_{-i}^h) \right],$$

where $\delta_i^h = \delta_i / [\delta_i + (1 - \delta_i) \pi_i(S_{-i}(h))]$ (the GBU update of δ_i) and π_i^h is the Bayesian update of π_i

whenever $\delta_i^h < 1$.

One-step deviations Consider a given a history $h \in H^{t-1}$ and a strategy profile $s \in S$. A *one-step deviation* in stage t by player i from her strategy s_i to the action $a_i \in A_i^h$ leads to the terminal history in Z^h determined by the continuation strategy profile $\langle a_i, s_i^h(-t), s_{-i}^h \rangle$, where $s^h \in S^h$, is the continuation of the strategy profile s starting in stage t from history h , and $s_i^h(-t)$ is player i 's component of that strategy profile except for her choice of action in stage t . This enables us to separate player i 's decision at stage t from the decisions of other players including her own past and future selves.

4 Equilibrium Concept: Consistent Planning Equilibrium Under Ambiguity (CP-EUA)

Our solution concept is an equilibrium in beliefs. Players choose pure (behavior) strategies, but have possibly ambiguous beliefs about the strategy choice of their opponent. Each player is required to choose at every decision node an action, which must be optimal with respect to their updated beliefs. When choosing an action a player treats his or her own future strategy as given. Consistency is achieved by requiring that the support of these beliefs is concentrated on the opponent's best replies. Thus it is a solution concept in the spirit of the agent normal form.

Definition 4.1 Fix a multi-stage game $\langle \{1, 2\}, \overline{H}, u_i, i = 1, 2 \rangle$. A Consistent Planning Equilibrium Under Ambiguity (CP-EUA) is a profile of capacities $\langle \nu_1, \nu_2 \rangle$ such that for each player $i = 1, 2$,

$$s_i \in \text{supp } \nu_{-i} \Rightarrow V_i^h \left(s_i^h | \nu_i^h \right) \geq V_i^h \left(\left(a_i, s_i^h(-t) \right) | \nu_i^h \right),$$

for every $a_i \in A_i^h$, every $h \in H^{t-1}$, and every $t = 1, \dots, T$.

Remark 4.1 If $|\text{supp } \nu_i| = 1$ for $i = 1, 2$ we say that the equilibrium is singleton. Otherwise we say that it is mixed. Singleton equilibria are analogous to pure strategy Nash equilibria.

Remark 4.2 A CP-EUA satisfies the one step deviation principle. No player may increase their conditional neo-expected payoff by changing their action in a single time period. We do not include a formal proof since the result is implied by the notion of sequential rationality embodied in the definition

of a CP-EUA. Owing to the failure of dynamic consistency, however, the one step deviation principle does not imply sequential rationality.

CP-EUA requires that the continuation strategy that player i is planning to play from history $h \in H^{t-1}$ is in the support of their opponent's beliefs ν_{-i} . Moreover, the only strategies in the support of their opponent's beliefs ν_{-i} are ones in which the action choice at history $h \in H^t$ is optimal for player i given their updated capacity ν_i^h . This rules out “incredible threats” in dynamic games. Thus our solution concept is an ambiguous analogue of sub-game perfection.¹⁸ Since we require beliefs to be in equilibrium in each subgame, an equilibrium at the initial node will imply optimal behavior of each player at each decision node. In particular, players will have a consistent plan in the sense of Siniscalchi (2011).

Mixed equilibria should be interpreted as equilibria in beliefs. To illustrate this consider a given player (she). We assume that she chooses pure actions and any randomizing is in the mind of her opponent. We require beliefs to be consistent with actual behavior in the sense that pure strategies in the support of the beliefs induce behavior strategies, which are best responses at any node where the given player has the move. The combination of neo-expected payoff preferences and GBU updating is not, in general, dynamically consistent. A consequence of this is that in a mixed equilibrium some of the *pure strategies*, which the given player's opponents believe she may play, are not necessarily optimal at all decision nodes. This arises because her preferences may change when they are updated. In particular equilibrium pure strategies will typically not be indifferent at the initial node. However at any node she will choose *actions* which are best responses. All behavior strategies in the support of her opponents' beliefs will be indifferent. These issues do not arise with pure equilibria.

The following result establishes that when players are neo-expected payoff maximizers, an equilibrium exists for any exogenously given degrees of ambiguity and ambiguity attitudes.

Proposition 4.1 *Let Γ be a multi-stage game with 2 neo-expected payoff maximizing players. Then Γ has at least one CP-EUA for any given parameters $\alpha_1, \alpha_2, \delta_1, \delta_2$, where $0 \leq \alpha_i \leq 1$, $0 < \delta_i \leq 1$, for $i = 1, 2$.*

¹⁸Recall that in a multi-stage game a new subgame starts after any given history h .

5 The Centipede Game

In this section, we apply our analysis to the centipede game. This is a two-player game with perfect information. One aim is to see whether strategic ambiguity can contribute to explaining observed behavior. In this section we shall assume that both players have neo-expected payoff preferences.¹⁹ The centipede game was introduced by Rosenthal (1981) and studied in laboratory experiments by McKelvey and Palfrey (1992). A survey of subsequent experimental research can be found in Krokow, Colman, and Pulford (2016).

5.1 The Game

The centipede game may be described as follows. There are two people, player 1 (she) and player 2 (he). Between them is a table which contains $2M$ one-pound coins and a single two-pound coin. They move alternately. At each move there are two actions available. The player whose move it is may either pick up the two-pound coin in which case the game ends; or (s)he may pick up two one-pound coins keep one and give the other to his/her opponent; in which case the game continues. In the final round there is a single two-pound coin and two one-pound coins remaining. Player 2, who has the move, may either pick up the two-pound coin, in which case the game ends and nobody gets the one-pound coins; or may pick up the two one-pound coins keep one and give the other to his opponent, in which case the opponent also gets the two pound coin and the game ends. We label an action that involves picking up two one-pound coins by r (right) and an action of picking up the two-pound coin by d (down). The diagram below shows the final four decision nodes.

A standard backward induction argument establishes that there is a unique iterated dominance equilibrium. At any node the player, whose move it is, picks up the 2-pound coin and ends the game. There are other Nash equilibria. However these only differ from the iterated dominance equilibrium off the equilibrium path.

¹⁹One might criticise these preferences on the grounds that they only allow the best and worst outcomes to be over-weighted but do not allow over-weighting of other outcomes. In many cases the worst outcome is death. However it is likely that individuals would also be concerned about other bad outcomes such as serious injury and/or large monetary losses. Thus in many cases individuals may over-weight a number of bad outcomes rather than just the very worst outcome. Despite this potential problem, we believe this model is suitable for application to strategic situations in particular the centipede game. Our reason is that this game has focal best and worst outcomes, that is, the high payoff at the end and the low payoff from stopping the game.

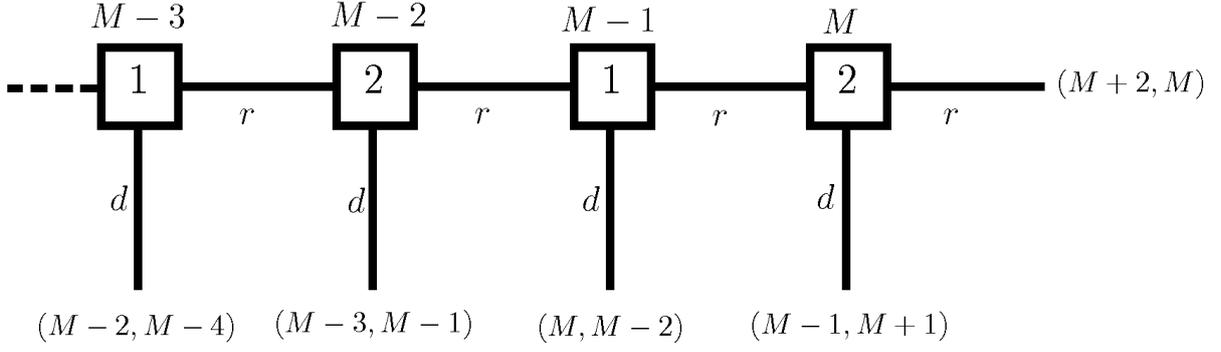


Figure 2: Last stages of the centipede game

5.2 Notation

Given the special structure of the centipede game, we can simplify our notation. The tree can be identified with the non-terminal nodes $H = \{1, \dots, M\}$. For simplicity we shall assume that M is an even number. The set of non-terminal nodes can be partitioned into the two player sets $H_1 = \{1, 3, \dots, M-1\}$ and $H_2 = \{2, 4, \dots, M\}$. It will be a maintained hypothesis that $M \geq 4$.

Strategies and Pay-offs A (pure) strategy for player i is a mapping $s_i : H_i \rightarrow \{r, d\}$. Given a strategy combination (s_1, s_2) set $m(s_1, s_2) := 0$ if d is never played, otherwise set $m(s_1, s_2) := m' \in H$ where m' is the first node where action d is played. The payoff of strategy combination (s_1, s_2) is: $u_1(s_1, s_2) = M + 2$, if $m(s_1, s_2) = 0$; $u_1(s_1, s_2) = m(s_1, s_2) + 1$ if $m(s_1, s_2)$ is odd and $u_1(s_1, s_2) = m(s_1, s_2) - 1$ otherwise. Likewise $u_2(s_1, s_2) = M$, if $m(s_1, s_2) = 0$; $u_2(s_1, s_2) = m(s_1, s_2) + 1$ if $m(s_1, s_2)$ is even; $u_2(s_1, s_2) = m(s_1, s_2) - 1$ otherwise. For each player $i = 1, 2$ and each node ρ in H_i , let s_i^ρ denote the threshold strategy defined as $s_i^\rho(m) = r$, for $m < \rho$; and $s_i^\rho(m) = d$ for $m \geq \rho$. Let s_i^∞ denote the strategy to play r always. Since threshold strategies s_i^ρ , $1 \leq \rho \leq M-1$, are weakly dominant they deserve special consideration. As we shall show, all equilibrium strategies are threshold strategies.

Subgames and Continuation Strategies For any node $m' \in H$ set $m(s_1, s_2 | m') := 0$ if d is not played at m' or thereafter. Otherwise set $m(s_1, s_2 | m') := m'' \in \{m', \dots, M\}$ where m'' is the first node where action d is played in the subgame starting at m' . We will write $u^i(s_1, s_2 | m')$ to denote the payoff of the continuation strategy from node m' , that is, $u^1(s_1, s_2 | m') = M + 2$ if $m(s_1, s_2 | m') = 0$; $u_1(s_1, s_2 | m') = m(s_1, s_2 | m') + 1$ if $m(s_1, s_2 | m')$ is odd; $u_1(s_1, s_2 | m') = m(s_1, s_2 | m') - 1$ otherwise. Similarly define $u_2(s_1, s_2 | m') = M$ if $m(s_1, s_2 | m') = 0$; $u_2(s_1, s_2 | m') = m(s_1, s_2 | m') + 1$ if $m(s_1, s_2 | m')$ is even; $u_2(s_1, s_2 | m') = m(s_1, s_2 | m') - 1$ otherwise.

5.3 Consistent Planning Equilibria Under Ambiguity

In this section we characterize the CP-EUA of the centipede game with symmetric neo-expected payoff maximizing players. That is, throughout this section we take Γ to be an M stage centipede game, where M is an even number no less than 4, and in which both players are neo-expected payoff maximizers with $\delta_1 = \delta_2 = \delta \in [0, 1]$ and $\alpha_1 = \alpha_2 = \alpha \in [0, 1]$.

There are three possibilities, cooperation continues until the final node, there is no cooperation at any node or there is a mixed equilibrium. As we shall show below, a mixed equilibrium also involves a substantial amount of cooperation. The first proposition shows that if there is sufficient ambiguity and players are sufficiently optimistic the equilibrium involves playing “right” until the final node. At the final node Player 2 chooses “down” since it is a dominant strategy.

Proposition 5.1 *For $\delta(1 - \alpha) \geq \frac{1}{3}$, there exists a CP-EUA $\langle \nu_1(\cdot|\alpha, \delta, \pi_1), \nu_2(\cdot|\delta, \alpha, \pi_2) \rangle$, with $\pi_1(s_2^*) = \pi_2(s_1^*) = 1$ for the strategy profile $\langle s_1^*, s_2^* \rangle$ in which $m(s_1^*, s_2^*) = M$. This equilibrium will be unique provided the inequality is strict.*

This confirms our intuition. Ambiguity-loving preferences can lead to cooperation in the centipede game. To understand this result, observe that $\delta(1 - \alpha)$ is the decision-weight on the best outcome in the Choquet integral. Cooperation does not require highly ambiguity loving preferences. A necessary condition for cooperation is that ambiguity-aversion is not too high i.e. $\alpha \leq \frac{2}{3}$. Such ambiguity-attitudes are not implausible, since Kilka and Weber (2001) experimentally estimate that $\alpha = \frac{1}{2}$.

Recall that players do not cooperate in the Nash equilibrium. We would expect that ambiguity-aversion makes cooperation less likely, since it increases the attractiveness of playing down which offers a low but ambiguity-free payoff. The next result finds that, provided players are sufficiently ambiguity-averse, non-cooperation at every node is an equilibrium.

Proposition 5.2 *For $\alpha \geq \frac{2}{3}$, there exists a CP-EUA $\langle \nu_1(\cdot|\alpha, \delta, \pi_1), \nu_2(\cdot|\delta, \alpha, \pi_2) \rangle$, with $\pi_1(s_2^*) = \pi_2(s_1^*) = 1$ for the strategy profile $\langle s_1^*, s_2^* \rangle$, in which $m(s_1^*, s_2^*|m') = m'$ at every node $m' \in H$. This equilibrium will be unique provided the inequality is strict.*

It is perhaps worth emphasizing that pessimism must be large in order to induce players to exit at every node. If $\frac{1}{2} < \alpha < \frac{2}{3}$ the players overweight bad outcomes more than they overweight good outcomes. However non-cooperation at every node is not an equilibrium in this case even though players are fairly pessimistic about their opponents' behavior.

We proceed to study the equilibria when $\alpha < \frac{2}{3}$ and $\delta(1 - \alpha) < \frac{1}{3}$. This case is interesting, since Kilka and Weber (2001) estimate parameter values for α and δ in a neighborhood of $\frac{1}{2}$. The next result shows that there is no singleton equilibrium for these parameter values and characterizes the mixed equilibria which arise. Interestingly, the equilibrium strategies imply continuation for most nodes. This supports our hypothesis that ambiguity-loving can help to sustain cooperation.

Proposition 5.3 *Assume that $\delta(1 - \alpha) < \frac{1}{3}$ and $\alpha < \frac{2}{3}$. Then:*

1. Γ does not have a singleton CP-EUA;

2. there exists a CP-EUA $\langle \nu_1(\cdot|\alpha, \delta, \pi_1), \nu_2(\cdot|\delta, \alpha, \pi_2) \rangle$ in which,

(a) player 1 believes with degree of ambiguity δ that player 2 will choose his strategies with

(ambiguous) probability $\pi_1(s_2) = p$ for $s_2 = s_2^M$; $1 - p$, for $s_2 = s_2^{M-2}$; $\pi_1(s_2) = 0$, otherwise, where $p = \frac{\delta(2-3\alpha)}{1-\delta}$;

(b) player 2 believes with degree of ambiguity δ that player 1 will choose her strategies with

(ambiguous) probability $\pi_2(s_1) = q$ for $s_1 = s_1^{M-1}$; $1 - q$, for $s_1 = s_1^{M-3}$; $\pi_2(s_1) = 0$, otherwise, where $q = \frac{1-3\delta(1-\alpha)}{3(1-\delta)}$;

(c) The game will end at $M - 2$ with player 2 exiting, at $M - 1$ with player 1 exiting, or at M with player 2 exiting.

Notice that for the profile of admissible capacities $\langle \nu_1(\cdot|\alpha, \delta, \pi_1), \nu_2(\cdot|\delta, \alpha, \pi_2) \rangle$ specified in Proposition 5.3 to constitute a CP-EUA, we require player 2's "theory" about the "randomization" of player 1's choice of action at node $M - 1$ should make player 2 at node $M - 2$ indifferent between selecting either d or r . That is, $M - 1 = (1 - \delta)((1 - q)(M - 2) + q(M + 1)) + \delta(\alpha(M - 2) + (1 - \alpha)(M + 1))$, which solving for q yields,

$$q = \frac{1 - 3\delta(1 - \alpha)}{3(1 - \delta)}. \quad (2)$$

This is essentially the usual reasoning employed to determine the equilibrium 'mix' with standard expected payoff maximizing players.

The situation for player 1 is different, however, since her perception of the "randomization" undertaken by player 2 over his choice of action at node $M - 2$ increases the ambiguity player 1 experiences at node $M - 1$. Given full Bayesian updating, this should generate enough ambiguity for player 1 so that she is indifferent between her two actions at node $M - 1$ given her "theory" that

Player 2 will choose d at node M . More precisely, given the GBU of Player 1's belief conditional on reaching node $M - 1$, player 1 should be indifferent between selecting either d or r ; that is, $M = (1 - \delta^{M-1}(1 - \alpha))(M - 1) + \delta^{M-1}(M + 2)$, where $\delta^{M-1} = \frac{\delta}{\delta + (1-\delta)p}$. Solving for p yields,

$$p = \frac{\delta(2 - 3\alpha)}{1 - \delta}. \quad (3)$$

Thus substituting p into the expression above for δ^{M-1} we obtain $\delta^{M-1} = \frac{1}{3(1-\alpha)}$ and $\delta^{M-1}(1 - \alpha) = \frac{1}{3}$, as required.

Remark 5.1 *It may at first seem puzzling that as $\delta \rightarrow 0$, we have $q \rightarrow \frac{1}{3}$, $p \rightarrow 0$ and $\delta^{M-1} = \frac{1}{3(1-\alpha)}$, for all $\delta \in \left(0, \frac{1}{3(1-\alpha)}\right)$, and yet for $\delta = 0$ (that is, with standard expected payoff maximizing players) by definition $\delta^{M-1} = 0$ and the unique equilibrium entails both players choosing d at every node, so in particular, $q = p = 0$. This discontinuity, is simply a consequence of the fact that (for fixed δ) $\frac{\delta}{\delta + (1-\delta)p} \rightarrow 1$ as $p \rightarrow 0$ in contrast to an intuition that the updated degree of ambiguity δ^{M-1} should converge to zero as $\delta \rightarrow 0$. Notice that for any (constant) $p > 0$, $\delta \rightarrow 0$ would indeed imply $\frac{\delta}{\delta + (1-\delta)p} \rightarrow 0$. However, to maintain an equilibrium of the type characterized in Proposition 5.3, p has to increase sufficiently fast to maintain $\delta^{M-1} = \frac{1}{3(1-\alpha)}$.*

The discontinuity at $\delta = 0$ is puzzling if the intuition is guided by what one knows about mixed strategies and exogenous randomizations of payoffs in perturbed games. Moreover this argues against interpreting any limit of a sequence of CP-EUA as $\delta \rightarrow 0$ as constituting a possible refinement of subgame perfect (Nash) equilibrium. Without optimism, there is no discontinuity, but then we are no longer able to explain the observed continuation in centipede games.²⁰

The mixed equilibria occur when $\alpha < \frac{2}{3}$ and $\delta(1 - \alpha) < \frac{1}{3}$. These parameter values could be described as situations of low ambiguity and low pessimism. On the equilibrium path players are not optimistic enough, given the low degrees of ambiguity, in order to play “right” at all nodes. However low pessimism makes them optimistic enough for playing “right” once they are off the equilibrium path whenever it is not a dominated strategy. This difference in behavior on and off the equilibrium path is the reason for non-existence of a singleton equilibrium.

In the mixed equilibrium the support of the original beliefs would contain two pure strategies, which player 1 has a strict preference between. However at any node where they differ the behavior

²⁰We thank stimulating comments and suggestions from David Levine and Larry Samuelson for motivating this remark.

strategies which they induce are indifferent. (In these circumstances player 2 might well experience ambiguity concerning which strategy player 1 is following.)

The conditions $\alpha \gtrless \frac{2}{3}$ and $\delta(1 - \alpha) \gtrless \frac{1}{3}$ characterize the parameter regions for the three types of *CP-EUA* equilibria. These are shown in figure 3. For strong pessimism ($\alpha > \frac{2}{3}$) players will always exit (red region), while for sufficient optimism and ambiguity ($\delta(1 - \alpha) > \frac{1}{3}$) players will always continue (blue region). Kilka and Weber (2001) experimentally estimate the parameters of the neo-additive model as $\delta = \alpha = \frac{1}{2}$. For parameters in a neighborhood of these values only the mixed equilibrium exists. This would be compatible with a substantial degree of cooperation.

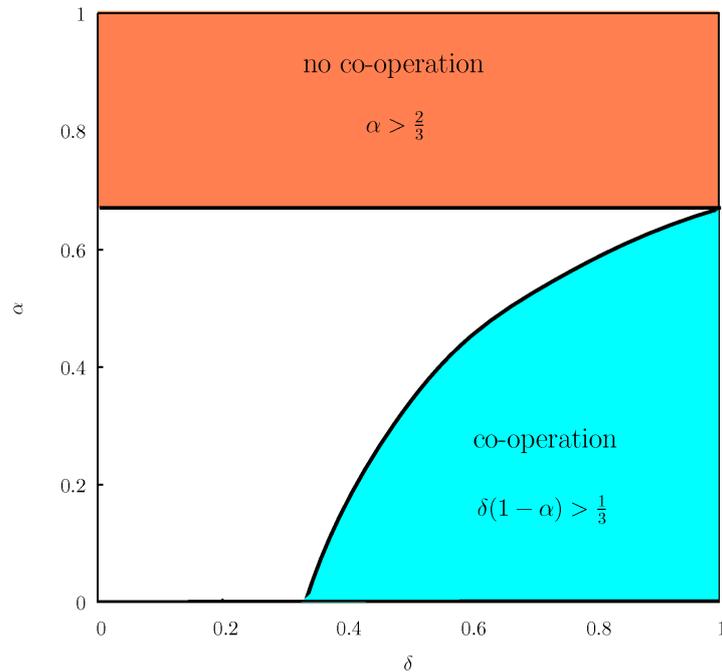


Figure 3: Equilibrium regions

6 Bargaining

The alternating offer bargaining game, was developed by Stahl (1972) and Rubinstein (1982), has become one of the most intensely studied models in economics, both theoretically and experimentally. In its shortest version, the ultimatum game, it provides a prime example for a subgame perfect Nash equilibrium prediction at odds with experimental behavior. The theoretical prediction is of an initial offer of the smallest possible share of a surplus (often zero) followed by acceptance. However experimental results show that the initial offers range around a third of the surplus which is often, but by far not always, accepted.

In bargaining games lasting for several rounds, the same subgame perfect equilibrium predicts a

minimal offer depending on the discount rate and the length of the game, which will be accepted in the first round. Experimental studies show, however, that players not only make larger offers than suggested by the equilibrium but also do not accept an offer in the first round (Roth (1995), p. 293). In a game of perfect information rational agents should not waste resources by delaying agreement.

In order to accommodate the observed delays, game-theoretic analysis has suggested incomplete information about the opponent's payoffs. Though it can be shown that incomplete information can lead players to reject an offer, the general objection to this explanation advanced in Forsythe, Kennan, and Sopher (1991) remains valid:

“In a series of recent papers, the Roth group has shown that even if an experiment is designed so that each bargainer knows his opponent's utility payoffs, the information structure is still incomplete. In fact, because we can never control the thoughts and beliefs of human subjects, it is impossible to run a complete information experiment. More generally, it is impossible to run an incomplete information experiment in which the experimenter knows the true information structure. Thus we must be willing to make conjectures about the beliefs which subjects might plausibly hold, and about how they may reasonably act in light of these beliefs. (p.243)”

In this paper we suggest another explanation. Following Luce and Raiffa (1957), p.275, we will assume that players view their opponent's behavior as ambiguous. Though this uncertainty will be reduced by their knowledge about the payoffs of the other player and their assumption that opponents will maximize their payoff, players cannot be completely certain about their prediction. As we will show such ambiguity can lead to delayed acceptance of offers.

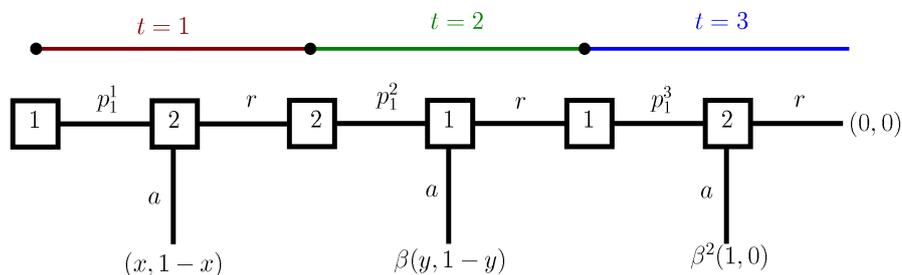


Figure 4: The bargaining game

Consider the bargaining game in figure 4. Without ambiguity, backward induction predicts a split of $\langle \beta(1 - \beta), 1 - \beta(1 - \beta) \rangle$ which will be accepted in period $t = 1$. Delay is not sensible because the best a player can expect from rejecting this offer is the same payoff (modulo the discount factor) a period later. Depending on the discount factor β the lion's share will go to the player who makes the

offer in the last stage when the game turns into an ultimatum game.

Suppose now that a player feels some ambiguity about such equilibrium behavior. Such ambiguity appears particularly reasonable because the incentives of the two players are delicately balanced. If a player has even a small degree of optimism, they may consider it possible that, by deviating from the expectations of the equilibrium path, the opponent may accept an offer which is more favorable for them. Hence, there may be an incentive to “test the water” by deviating from the equilibrium path. Note that this may be a low-cost deviation since, by returning to the previous path, just the discount is lost. Hence, if the discount is low, that is., the discount factor β is high, a small degree of optimism may suffice.

Decision makers with neo-expected payoff preferences who face ambiguity $\delta > 0$ and update their beliefs according to the GBU rule give some extra weight $(1 - \alpha)$ to the best expected payoff and α to the worst expected payoff and update their beliefs to complete ambiguity, $\delta = 1$, if an event occurs which has probability zero according to their focal (additive) belief π . Hence, off the equilibrium path updates are well-defined but result in complete uncertainty. A decision maker with neo-additive beliefs will evaluate their strategies following an out-of-equilibrium move and, therefore probability zero event of π by their best and worst outcomes. Hence, from an optimistic perspective, asking for a high share may have a chance of being accepted resulting in some expected gain which can be balanced against the loss of discount associated with a rejection. Whether a strategy resulting in a delay is optimal will depend on the degree of ambiguity δ , the degree of optimism $1 - \alpha$ and the discount factor β .

The following result supports this intuition. With ambiguity and some optimism, delayed agreement along the equilibrium path may occur in a CP-EUA equilibrium.

Proposition 6.1 *If $\frac{\alpha - (1 - \alpha)\beta}{1 - (1 - \alpha)[\max\{1 - (1 - \alpha)\beta, \beta\} + \beta]} \geq \delta$, then there exists a CP-EUA*

$\langle (\delta, \alpha, \pi_1), (\delta, \alpha, \pi_2) \rangle$ such that $\pi_1(s_2^) = \pi_2(s_1^*) = 1$ for the following strategy profile (s_1^*, s_2^*) :*

- *at $t = 1$, player 1 proposes division $\langle x^*, 1 - x^* \rangle = \langle 1, 0 \rangle$, player 2 accepts a proposed division $\langle x, 1 - x \rangle$ if and only if $x \leq 1 - [(1 - \delta)(1 - (1 - \alpha)\beta) + \delta(1 - \alpha)\beta \max\{1 - (1 - \alpha)\beta, \beta\}]$;*
- *at $t = 2$,*
 - *if player 1's proposed division in $t = 1$ was $\langle x, 1 - x \rangle = \langle 1, 0 \rangle$, then player 2 proposes division $\langle y^*, 1 - y^* \rangle = \langle (1 - \alpha)\beta, 1 - (1 - \alpha)\beta \rangle$ and Player 1 accepts a proposed division $\langle y, 1 - y \rangle$ if and only if $y \geq (1 - \alpha)\beta$;*

- otherwise, Player 2 proposes division $\langle \hat{y}, 1 - \hat{y} \rangle = \langle 0, 1 \rangle$, Player 1 accepts the proposed division $\langle y, 1 - y \rangle$ if and only if $y \geq (1 - \alpha) \beta$;
- at $t = 3$, player 1 proposes division $\langle z^*, 1 - z^* \rangle = \langle 1, 0 \rangle$ and player 2 accepts any proposed division $\langle z, 1 - z \rangle$.²¹

7 Relation to the Literature

This section relates the present paper to the existing literature. First we consider our own previous research followed by the relation to other theoretical research in the area. Finally we discuss the experimental evidence.

7.1 Ambiguity in Games

Most of our previous research has considered normal form games e.g. Eichberger and Kelsey (2014). The present paper extends this by expanding the class of games. Two earlier papers study a limited class of extensive form games, Eichberger and Kelsey (2004) and Eichberger and Kelsey (1999). These focus on signalling games in which each player only moves once. Consequently dynamic consistency is not a major problem. Signalling games may be seen as multi-stage games with only two stages and incomplete information. The present paper relaxes this restriction on the number of stages but has assumed complete information. The price of increasing the number of stages is that we are forced to consider dynamic consistency.

Rothe (2011) models each player's belief in an extensive form game as the special subclass of a neo-additive capacity with extreme pessimism (that is, $\alpha_i = 1$). However, he interprets the weight on the additive part of this capacity as the probability the player thinks her opponent is rational with the complementary weight being the probability she thinks her opponent is *irrational*. Furthermore, these (conditional) weights are specified *exogeneously* for each decision node of the player. It can thus be viewed as a Choquet expected utility extension of the irrational type literature discussed in section 8.1 below.

Hanany, Klibanoff, and Mukerji (2016) (henceforth HKM) also present a theory of ambiguity in multi-stage games. However they have made a number of different modelling choices. Firstly they consider games of incomplete information. In their model, there is ambiguity concerning the type of

²¹Note this is irrespective of Player 1's proposed division $\langle x, 1 - x \rangle$ at $t = 1$, and irrespective of Player 2's proposed division $\langle y, 1 - y \rangle$ at $t = 2$:

the opponent while their strategy is unambiguous. In contrast in our theory there is no type space and we focus on strategic ambiguity. However we believe that there is not a vast difference between strategic ambiguity and ambiguity over types. It would be straightforward to add a type space to our model, while HKM argue that strategic uncertainty can arise as a reduced form of a model with type uncertainty. Other differences are that HKM represent ambiguity by the smooth model, they use a different rule for updating beliefs and strengthen consistent planning to dynamic consistency. A cost of this is that they need to adopt a non-consequentialist decision rule. Thus current decisions may be affected by options which are no longer available.

We conjecture that similar results could have been obtained using the smooth model. However, the GBU rule has the advantage that it defines beliefs both on and off the equilibrium path. In contrast, with the smooth rule, beliefs off the equilibrium path are to some extent arbitrary. In addition, we note that since there is little evidence that individuals are dynamically consistent, this assumption is more suitable for a normative model rather than a descriptive theory. As HKM show, dynamic consistency imposes strong restrictions on preferences and how they are updated.

Jehiel (2005) proposes a solution concept which he refers to as *analogy-based equilibrium*. In this a player identifies similar situations and forms a single belief about their opponent's behavior in all of them. These beliefs are required to be correct in equilibrium. For instance in the centipede game a player might consider their opponent's behavior at all the non-terminal nodes to be analogous. Thus they may correctly believe that the opponent will play right with high probability at the average node, which increases their own incentive to play right. (The opponent perceives the situation similarly.) Jehiel predicts that either there is no cooperation or cooperation continues until the last decision node. This is not unlike our own predictions based on ambiguity.

What is common between his theory and ours is that there is an "averaging" over different decision nodes. In his theory this occurs through the perceived analogy classes, while in ours averaging occurs via the decision-weights in the Choquet integral. We believe that an advantage of our approach is that the preferences we consider have been derived axiomatically and hence are linked to a wider literature on decision theory.

7.2 Experimental Papers

Our paper predicts that ambiguity about the opponent's behavior may significantly increase cooperation above the Nash equilibrium level in the centipede game. This prediction is broadly confirmed

by the available experimental evidence, (for a survey see Krokow, Colman, and Pulford (2016)).

McKelvey and Palfrey (1992) study 4 and 6-stage centipede games with exponential payoffs. They find that most players play right until the last 3-4 stages, after which cooperation appears to break down randomly. This is compatible with our results on the centipede game which predict that cooperation continues until near the end of the game.²²

Our paper makes the prediction that either there will be no cooperation in the centipede game or that cooperation will continue until the last three stages. In the latter case it will either break down randomly in a mixed equilibrium or break down at the final stage in a singleton equilibrium. In particular the paper predicts that cooperation will not break down in the middle of a long centipede game. This can in principle be experimentally tested. However we would note that it is not really possible to refute our predictions in a 4-stage centipede as used by McKelvey and Palfrey (1992) . Thus there is scope for further experimental research on longer games.²³

8 Conclusion

This paper has studied extensive form games with ambiguity. This is done by constructing a thought experiment, where we introduce ambiguity but otherwise make as few changes to standard models as possible. We have proposed a solution concept for multi-stage games with ambiguity. An implication of this is that singleton equilibria may not exist in games of complete and perfect information. This is also demonstrated by the fact that the centipede game only has mixed equilibria for some parameter values. We have shown that ambiguity-loving behavior may explain apparently counter-intuitive properties of Nash equilibrium in the Centipede game and non-cooperative bargaining. It also produces predictions closer to the available evidence than Nash equilibrium.

8.1 Irrational Types

As mentioned in the introduction, economists have been puzzled about the deviations from Nash predictions in a number of games such as the centipede game, the repeated prisoners' dilemma and the chain store paradox. In the present paper we have attempted to explain this behavior as a re-

²²In an earlier draft we proved that there existed regions of the parameter space in which results analogous to Propositions 5.2 and 5.3 held for players with exponential utility.

²³There are a number of other experimental papers on the centipede game. However many of them do not study the version of the game presented in this paper. For instance they may consider a constant sum centipede or study the normal form. It is not clear that our predictions will apply to these games. Because of this, we do not consider them in this review. For a survey see Krokow, Colman, and Pulford (2016).

sponse to ambiguity. Previously it has been common to explain these deviations by the introduction of an “irrational type” of a player. This converts the original game into a game of incomplete information where players take into consideration a small probability of meeting an irrational opponent. An “irrational” player is a type whose payoffs differ from the corresponding player’s payoffs in the original game. In such modified games of incomplete information, it can be shown that the optimal strategy of a “rational” player may involve imitating the “irrational” player in order to induce more favorable behavior by his/her opponents. This method is used to rationalize observed behavior in the repeated prisoner’s dilemma, (Kreps, Milgrom, Roberts, and Wilson (1982)), and in the centipede game (McKelvey and Palfrey (1992)).

There are at least two reasons why resolving the conflict between backward induction and observed behavior by introducing “irrational” players may not be the complete answer. First, games of incomplete information with “irrational” players predict with small probabilities that two irrational types will confront each other. Hence, this should appear in the experimental data. Secondly, in order to introduce the appropriate “irrational” types, one needs to know the observed deviations from equilibrium behavior. Almost any type of behavior can be justified as a response to some kind of irrational opponent. It is plausible that one’s opponent may play tit for tat in the repeated prisoners’ dilemma. Thus an intuitive account of cooperation in the repeated prisoners’ dilemma may be based on a small probability of facing an opponent of this type. However, for most games, there is no such focal strategy which one can postulate for an irrational type to adopt. Theory does not help to determine which irrational types should be considered and hence does not make usually clear predictions. In contrast our approach is based on axiomatic decision theory and can be applied to any multi-stage game.

8.2 Directions for Future Research

In the present paper we have focused on multi-stage games. There appears to be scope for extending our analysis to a larger class of games. For instance, we believe that it would be straightforward to add incomplete information by including a type space for each player. Extensions to multi-player games are possible. If there are three or more players it is usual to assume that each one believes that his/her opponents act independently. At present it is not clear as to how one should best model independence of ambiguous beliefs.²⁴

²⁴There are still some differences of opinion among the authors of this paper on this point.

It should also be possible to extend the results to a larger class of preferences. Our approach is suitable for any ambiguity model which maintains a separation between beliefs and tastes and allows a suitable support notion to be defined. In particular both the multiple priors and smooth models of ambiguity fit these criteria. These models represent beliefs by a set of probabilities. A suitable support notion can be defined in terms of the intersection of the supports of the probabilities in this set of beliefs. This is the inner support of Ryan (2002).

A natural application is to finitely repeated games. Such games have some features in common with the centipede game. If there is a unique Nash equilibrium then backward induction implies that there is no scope for cooperation in the repeated game. However in examples, such as the repeated prisoners' dilemma, intuition suggests that some cooperation should be possible.

We believe the model is suitable for applications in financial markets. In particular phenomena such as asset price bubbles and herding have some features in common with the centipede game. Thus we believe that our analysis could be used to study them. In an asset price bubble the value of a security rises above the level, which can be justified by fundamentals. Individuals continue buying even though they know the price is too high since they believe it will rise still further. Thus at every step a buyer is influenced by the perception that the asset price will continue to rise even though they are aware it cannot rise for ever. This is somewhat similar to observed behavior in the centipede game, where players choose "right" many times even though they know that cooperation cannot last indefinitely. Reasoning analogous to that of the present paper may be useful to explain an asset price bubble in terms of ambiguity-loving behavior.

References

- AZRIELI, Y., AND R. TEPER (2011): "Uncertainty aversion and equilibrium existence in games with incomplete information," *Games and Economic Behavior*, 73(2), 310–317.
- BAUMEISTER, R. F., AND K. D. VOHS (eds.) (2004): *Handbook of Self-regulation: Research, Theory and Applications*. Guilford Publications.
- BOSE, S., AND L. RENO (2014): "Mechanism Design with Ambiguous Communication Devices," *Econometrica*, 82(5), 1853–1872.
- 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(1), 538–567.

- DOW, J., AND S. R. C. WERLANG (1994): “Nash Equilibrium under Uncertainty: Breaking Down Backward Induction,” *Journal of Economic Theory*, 64, 305–324.
- 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.
- (2012): “When is Ambiguity-Attitude Constant?,” *Journal of Risk and Uncertainty*, 45, 239–263.
- (2016): “Randomisation and Dynamic Consistency,” *Economic Theory*, 62, 547–566.
- EICHBERGER, J., AND D. KELSEY (1999): “Education Signalling and Uncertainty,” in *Beliefs Interactions and Preferences*, ed. by Nau, Grøn, and Machina. Kluwer.
- (2000): “Non-Additive Beliefs and Strategic Equilibria,” *Games and Economic Behavior*, 30, 183–215.
- (2004): “Sequential Two-Player Games with Ambiguity,” *International Economic Review*, 45, 1229–1261.
- (2014): “Optimism and Pessimism in Games,” *International Economic Review*, 55, 483–505.
- EISENBERG, N., C. L. SMITH, AND T. L. SPINRAD (2004): “Effortful Control. Relations with Emotion Regulation, Adjustment, and Socialization in Childhood,” in *Handbook of Self-regulation: Research, Theory and Applications.*, ed. by R. F. Baumeister, and K. D. Vohs, chap. Chapter 14:, pp. 263–283. Guilford Publications.
- EPSTEIN, L. G., AND M. SCHNEIDER (2003): “Recursive Multiple-Priors,” *Journal of Economic Theory*, 113, 1–31.
- FORSYTHE, R., J. KENNAN, AND B. SOPHER (1991): “Dividing a Shrinking Pie: An Experimental Study of Strikes in Bargaining Games with Complete Information,” *Research in Experimental Economics*, 4, 223–268.
- 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 a 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.
- GRANT, S., I. MENEGHEL, AND R. TOURKY (2016): “Savage games,” *Theoretical Economics*, 11, 641–682.

- GREINER, B. (2016): “Strategic Uncertainty Aversion in Bargaining: Experimental Evidence,” *working paper, University of New South Wales*.
- GÜTH, W., R. SCHMITTBERGER, AND B. SCHWARZE (1982): “An Experimental Analysis of Ultimatum Bargaining,” *Journal of Economic Behavior and Organization*, 3, 367–388.
- HANANY, E., P. KLIBANOFF, AND S. MUKERJI (2016): “Incomplete information games with ambiguity averse players,” *working paper, Northwestern University*.
- HORIE, M. (2013): “Re-Examination on Updating Choquet Beliefs,” *Journal of Mathematical Economics*, 49, 467–470.
- HURWICZ, L. (1951): “Optimality Criteria for Decision Making under Ignorance,” Discussion paper, Cowles Commission.
- JEHIEL, P. (2005): “Analogy-based expectation equilibrium,” *Journal of Economic Theory*, 123, 81–104.
- KAHNEMAN, D., AND A. TVERSKY (1979): “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, 47, 263–291.
- KAJII, A., AND T. UI (2005): “Incomplete Information Games with Multiple Priors,” *The Japanese Economic Review*, 56(3), 332–351.
- KARNI, E., AND Z. SAFRA (1989): “Ascending bid auctions with behaviorally consistent bidders,” *Annals of Operations Research*, 19, 435–446.
- KELLNER, C., AND M. T. LEQUEMENT (2015): “Endogenous ambiguity in cheap talk,” *working paper, University of Bonn*.
- KILKA, M., AND M. WEBER (2001): “What Determines the Shape of the Probability Weighting Function under Uncertainty?,” *Management Science*, 47, 1712–1726.
- KREPS, D., P. MILGROM, J. ROBERTS, AND R. WILSON (1982): “Rational Cooperation in the Finitely Repeated Prisoner’s Dilemma,” *Journal of Economic Theory*, 27, 253–279.
- KROKOW, E. M., A. M. COLMAN, AND B. PULFORD (2016): “Cooperation in repeated interactions: A systematic review of Centipede game experiments, 1992-2016,” *European Review of Social Psychology*, 27, 231–282.
- LO, K. C. (1999): “Extensive Form Games with Uncertainty-Averse Players,” *Games and Economic Behavior*, 28, 256–270.
- LUCE, R. D., AND H. RAIFFA (1957): *Games and Decisions*. Dover Books on Mathematics.
- MARINACCI, M. (2000): “Ambiguous Games,” *Games and Economic Behavior*, 31, 191–219.
- MCKELVEY, R. D., AND T. R. PALFREY (1992): “An Experimental Study of the Centipede Game,” *Econometrica*, 60, 803–836.

- NEUMANN, J. V., AND O. MORGENSTERN (1944): *The Theory of Games and Economic Behavior*. Princeton University Press, New Jersey.
- OSBORNE, M. J., AND A. RUBINSTEIN (1994): *A course in game theory*. MIT Press.
- QUIGGIN, J. (1982): "A Theory of Anticipated Utility," *Journal of Economic Behavior and Organization*, 3, 323–334.
- RIEDEL, F., AND L. SASS (2014): "Ellsberg games," *Theory and Decision*, 76, 469–509.
- ROSENTHAL, R. (1981): "Games of Perfect Information, Predatory Pricing, and the Chain Store," *Journal of Economic Theory*, 25(1), 92–100.
- ROTH, A. E. (1995): "Bargaining Experiments," in *Handbook of Experimental Economics*, ed. by J. H. Kagel, and A. E. Roth, chap. 4, pp. 253–348. Princeton University Press.
- ROTHER, J. (2011): "Uncertainty aversion and equilibrium in extensive games," in *Contributions to game theory and management*, ed. by L. A. Petrosyan, and N. A. Zenevich, vol. IV of *The Fourth International Conference Game Theory and Management*, pp. 389–406, St Petersburg, Russia. St Petersburg University Press.
- RUBINSTEIN, A. (1982): "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50(1), 97–109.
- RYAN, M. J. (2002): "What Do Uncertainty-Averse Decision-Makers Believe?," *Economic Theory*, 20, 47–65.
- SARIN, R., AND P. WAKKER (1998): "Revealed Likelihood and Knightian Uncertainty," *Journal of Risk and Uncertainty*, 16, 223–250.
- SCHMEIDLER, D. (1989): "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57, 571–587.
- SELTEN, R. (1978): "The chain store paradox," *Theory and Decision*, 9, 127–159.
- SINISCALCHI, M. (2011): "Dynamic Choice Under Ambiguity," *Theoretical Economics*, 6, 379–421.
- SKIADAS, C. (1997): "Conditioning and aggregation of preferences," *Econometrica*, 65, 347–367.
- STAHL, I. (1972): *Bargaining Theory*, Stockholm School of Economics. Stockholm.
- STROTZ, R. (1955): "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 23, 165–180.
- TRAUTMANN, S., AND G. DE KUILEN (2015): "Ambiguity Attitudes," in *The Wiley Blackwell Handbook of Judgement and Decision Making*, ed. by G. Keren, and G. Wu. Wiley.

A Appendix: Proofs

A.1 Existence of Equilibrium

In this sub-appendix we present the proof of the existence result, Proposition 4.1.

The strategy of our proof is to associate with Γ a modified game Γ' , which is based on the agent-normal form of Γ . We show that Γ' has a Nash equilibrium and then use this to construct a CP-EUA for Γ . The game Γ' has 2θ players.²⁵ A typical player is denoted by $i_{h(t)}$, $h(t) \in H \setminus Z$, $i = 1, 2$. Thus there are 2 players for each decision node in Γ .

The strategy set of Player $i_{h(t)}$, $\Sigma_{i_{h(t)}} = \Delta(A_i^{h(t)})$ is the set of all probability distributions over $A_i^{h(t)}$, with generic element $s_{i_{h(t)}} \in \Sigma_{i_{h(t)}}$, for $i = 1, 2$. Hence in game Γ' , Player $i_{h(t)}$ may choose any mixed strategy over the set of actions $A_i^{h(t)}$. Let $\pi(h(t), \rho)$ denote the probability of history $h(t)$ when the strategy profile is ρ . This is calculated according to the usual rules for reducing compound lotteries to simple lotteries. We shall suppress the arguments and write $\pi = \pi(h(t), \rho)$ when convenient. Let π_i denote the marginal of π on S_{-i} .

The payoff of Player $i_{h(t)}$ is $\phi_{i_{h(t)}} : \Sigma_{i_{h(t)}} \rightarrow \mathbb{R}$, defined by $\phi_{i_{h(t)}}(a_i^{h(t)}, s_i^{t+1}, s_{-i}^t) = \int u_i(a_i^{h(t)}, s_i^{t+1}, s_{-i}^t) d\nu_{i_{h(t)}}$, where ν_i is the neo-additive capacity on S_{-i} defined by $\nu_i(\emptyset) = 0$; $\nu_i(A) = \delta_i(1 - \alpha_i)\pi_i(A)$, $\emptyset \subsetneq A \subsetneq S_{-i}$; $\nu_i(S_{-i}) = 1$ and $\nu_{i_{h(t)}}$ is the GBU update of ν_i conditional on $h(t)$. Since $\nu_{i_{h(t)}}$ is neo-additive:

$$\begin{aligned} \phi_{i_{h(t)}}(a_i^{h(t)}, s_i^{t+1}, s_{-i}^t) &= \delta_{i_{h(t)}}(1 - \alpha_i) M_{i_{h(t)}}(a_i^{h(t)}, s_i^{t+1}, s_{-i}^t) + \delta_{i_{h(t)}}\alpha_i m_{i_{h(t)}}(a_i^{h(t)}, s_i^{t+1}, s_{-i}^t) \\ &\quad + (1 - \delta_{i_{h(t)}}) \mathbf{E}_{\pi_{h(t)}^i} u_i(h(t), a_i^{h(t)}, s_i^{t+1}, s_{-i}^t), \end{aligned} \quad (4)$$

where $M_{i_{h(t)}}(a_i^{h(t)}, s_i^{t+1}, s_{-i}^t) = \max_{s_{-i}^t \in S_{-i}^t} u_i(h^t, a_i^{h(t)}, s_i^{t+1}, s_{-i}^t)$, and $m_{i_{h(t)}}(a_i^{h(t)}, s_i^{t+1}, s_{-i}^t) = \min_{s_{-i}^t \in S_{-i}^t} u_i(h^t, a_i^{h(t)}, s_i^{t+1}, s_{-i}^t)$. Here $\pi_{h(t)}^i$ denotes the Bayesian update of π^i , given that node $h(t)$ has been reached and $h(t)$ has positive probability. (If $h(t)$ has probability 0, then $\delta_{i_{h(t)}} = 1$ and $\pi_{h(t)}^i$ can be any probability distribution over S_{-i}^t .)

If Player $i_{h(t)}$ plays a mixed strategy then his/her action may be described by a probability distribution ρ over $A_i^{h(t)}$, which is treated as an *ex-ante* randomization. Eichberger, Grant, and Kelsey

²⁵Recall $\theta = |H \setminus Z|$ denotes the number of non-terminal histories.

(2016) show that individuals will be indifferent to ex-ante randomizations. Hence it is evaluated as

$$\sum_{a \in A_i^{h(t)}} \rho(a) \phi_{i_{h(t)}}(a, s_i^{t+1}, s_{-i}^t). \quad (5)$$

It follows that $i_{h(t)}$'s preferences are linear and hence quasi-concave in his/her own strategy.²⁶

Likewise if one of $i_{h(t)}$'s "future selves" randomizes this is evaluated as

$\sum_{s_i^{t+1} \in S_i^{t+1}} \xi(a) \phi_{i_{h(t)}}(a, s_i^{t+1}, s_{-i}^t)$, where ξ is the probability distribution over S_i^{t+1} induced by future randomizations. This is treated as an ex-ante randomization because it is resolved before the strategic ambiguity arising from the choice of i 's opponent in the relevant subgame. We do not need to specify $i_{h(t)}$'s reaction to randomizations by his/her past selves since these are, by definition, already resolved at the point where the decision is made.

Lemma A.1 *The function $\phi_{i_{h(t)}}(a_i^{h(t)}, s_i^{t+1}, s_{-i}^t)$ is continuous in s , provided $1 \geq \delta_i > 0$, for $i = 1, 2$.*

Proof. Consider equation (4). First note that $\phi_{i_{h(t)}}$ depends directly on s via the $(a_i^{h(t)}, s_i^{t+1}, s_{-i}^t)$ term. It also depends indirectly on s since the degree of ambiguity $\delta_{i_{h(t)}}$ and $\pi_{i_{h(t)}}$ are functions of s .

It follows from our assumptions that the direct relation between s and ϕ is continuous. By equation (5), ϕ is continuous in $\pi_{i_{h(t)}}$. Thus we only need to consider whether $\delta_{i_{h(t)}}$ and $\pi_{i_{h(t)}}$ are continuous in s . Recall that $\pi_{i_{h(t)}}$ is the probability distribution over terminal nodes induced by the continuation strategies s_i^t, s_{-i}^t . Since this is obtained by applying the law of compound lotteries it depends continuously on s . By definition $\delta_i^h = \frac{\delta_i}{\delta_i + (1 - \delta_i)\pi_i(S_{-i}(h))}$. This is continuous in $\pi(h(t))$ provided the denominator is not zero, which is ensured by the condition $\delta_i > 0$. Since $\pi(h(t))$ is a continuous function of s , the result follows. ■

The next result establishes that the associated game Γ' has a standard Nash equilibrium.

Lemma A.2 *The associated game Γ' has a Nash equilibrium provided $1 \geq \delta_i > 0$, for $i = 1, 2$.*

Proof. In the associated game Γ' , the strategy set of a typical player, $i_{h(t)}$, is the set of all probability distributions over the finite set $A_i^{h(t)}$ and is thus compact and convex. By equation (5) the payoff, $\phi_{i_{h(t)}}$, of Player $i_{h(t)}$ is continuous in the strategy profile σ . Moreover $\phi_{i_{h(t)}}$ is quasi concave in own strategy by equation (5). It follows that Γ' satisfies the conditions of Nash's theorem and therefore has a Nash equilibrium in mixed strategies. ■

²⁶To clarify these remarks about randomization apply to the modified game Γ' . In the original game Γ there is an equilibrium in beliefs and no randomization is used.

Proposition 4.1 *Let Γ be a 2-player multi-stage game. Then Γ has at least one CP-EUA for any given parameters $\alpha_1, \alpha_2, \delta_1, \delta_2$, where $1 \leq \alpha_i \leq 0, 0 < \delta_i \leq 1$, for $i = 1, 2$.*

Proof. Let $\rho = \langle \rho^{i_{h(t)}} : i = 1, 2, h \in H \setminus Z \rangle$ denote a Nash equilibrium of Γ' . We shall construct a CP-EUA $\hat{\sigma}$ of Γ based on ρ . Note that ρ may be viewed as a profile of behavior strategies in Γ . Let \hat{s} denote the profile of mixed strategies in Γ , which corresponds to ρ , and let π denote the probability distribution which ρ induces over S . (If ρ is an equilibrium in pure strategies then π will be degenerate.) The beliefs of player i in profile $\hat{\sigma}$ are represented by a neo-additive capacity $\hat{\nu}_i$ on S_{-i} , defined by $\hat{\nu}_i(B) = \delta_i(1 - \alpha_i) + (1 - \delta_i)\pi_i(B)$, where $B \subseteq S_{-i}$ and π_i denotes the marginal of π on S_{-i} .

Let $\hat{\nu}$ denote the profile of beliefs $\hat{\nu} = \langle \hat{\nu}_1, \hat{\nu}_2 \rangle$. We assert that $\hat{\nu}$ is a CP-EUA of Γ . By Remark 4.2, it is sufficient to show that no player can increase his or her current utility by a one-step deviation. Consider a typical player j . Let $\hat{t}, 0 \leq \hat{t} \leq T$, be an arbitrary time period and consider a given history $\hat{h}(\hat{t})$ at time \hat{t} . Let $\tilde{a}_j^{\hat{t}} \in A_{-j}^{h(\hat{t})}$ denote an arbitrary action for j at history $\hat{h}(\hat{t})$. Since ρ is an equilibrium of Γ' ,

$$\phi_{j_{\hat{h}(\hat{t})}}(\tilde{a}_j^{\hat{t}}, s_j^{\hat{t}+1}, s_{-j}^{\hat{t}}) \geq \phi_{j_{\hat{h}(\hat{t})}}(\hat{a}_j^{\hat{t}}, s_j^{\hat{t}+1}, s_{-j}^{\hat{t}})$$

for any $\tilde{a}_j^{\hat{t}} \in \text{supp } \rho(j) = \text{supp } \hat{\nu}_j^{\hat{h}(\hat{t})}$, where $\rho(j)$ denotes the marginal of ρ on $A_j^{h(\hat{t})}$. Without loss of generality we may assume that $\tilde{a}_j^{\hat{t}} = \hat{s}_j^{\hat{h}(\hat{t})} \in \text{supp } \hat{\nu}_j$. By definition $\phi_{j_{\hat{h}(\hat{t})}}(\tilde{a}_j^{\hat{t}}, s_j^{\hat{t}+1}, s_{-j}^{\hat{t}}) = \int u_j \left(\hat{s}_j^{\hat{h}(\hat{t})}, \hat{s}_{\langle j, 1 \rangle}^{\hat{t}+1}, s_{-j}^{\hat{t}} \right) d\nu_j^{\hat{h}(\hat{t})}$, where $\nu_j^{\hat{h}(\hat{t})}$ is the GBU update of ν_j conditional on $\hat{h}(\hat{t})$. Since the behavior strategy $\hat{s}_j^{\hat{h}(\hat{t})}$ is by construction a best response at $\hat{h}(\hat{t})$, $\int u_j \left(\hat{s}_j^{\hat{h}(\hat{t})}, \hat{s}_j^{\hat{t}+1}, s_{-j}^{\hat{t}} \right) d\nu_j^{\hat{h}(\hat{t})} \geq \int u_j \left(\tilde{a}_j^{\hat{t}}, \hat{s}_j^{\hat{t}+1}, s_{-j}^{\hat{t}} \right) d\nu_j^{\hat{h}(\hat{t})}$, which establishes that $\hat{s}_j^{\hat{h}(\hat{t})}$ yields a higher payoff than the one step deviation to $\tilde{a}_j^{\hat{t}}$. Since both j and $\hat{h}(\hat{t})$ were chosen arbitrarily this establishes that it is not possible to improve upon $\hat{\sigma}$ by a one-step deviation. Hence by Remark 4.2, $\hat{\sigma}$ is a CP-EUA of Γ . \blacksquare

A.2 The Centipede Game

Proof of Proposition 5.1 We shall proceed by (backward) induction. The final node is M . At this node 2 plays d_M , which is a dominant strategy. This yields payoffs $\langle M - 1, M + 1 \rangle$.

Node $M - 1$ Now consider Player 1's decision at node $M - 1$. Assume that this node is on the equilibrium path.²⁷ The (Choquet) expected value of her payoffs are:

$$V_1^{M-1}(d_{M-1} | \nu_1(\cdot | \alpha, \delta, \pi_1)) = M,$$

²⁷This will be proved once we have completed the induction.

$$V_1^{M-1}(r_{M-1}|\nu_1(\cdot|\alpha, \delta, \pi_1)) = \delta(1-\alpha)(M+2) + \delta\alpha(M-1) + (1-\delta)(M-1) = M-1+3\delta(1-\alpha).$$

Thus r_{M-1} is a best response provided $\delta(1-\alpha) \geq \frac{1}{3}$. To complete the proof we need to show that $s_i^*(\rho) = r_\rho$ is the preferred action for all $\rho \in H_i$ with $i = 1, 2$.

Inductive step Consider node ρ . Assume ρ is on the equilibrium path. We make the inductive hypothesis that r_κ is a best response at all nodes κ , $\rho < \kappa < M-1$. There are two cases to consider:

Case 1 $\rho = 2\tau + 1$ Player 1 has the move. The expected value of her payoffs are:

$$V_1^\rho(d_\rho|\nu_1(\cdot|\alpha, \delta, \pi_1)) = \rho + 1, \quad V_1^\rho(r_\rho|\nu_1(\cdot|\alpha, \delta, \pi_1)) = \delta(1-\alpha)(M+2) + \delta\alpha\rho + (1-\delta)(M-1).$$

Thus r_ρ is a best response provided $(1-\delta\alpha)M + 2\delta(1-\alpha) - (1-\delta) \geq (1-\delta\alpha)\rho + 1$,

$\Leftrightarrow (1-\delta\alpha)(M-\rho) \geq 2-3\delta+2\delta\alpha$. Now $(1-\delta\alpha)(M-\rho) \geq 3(1-\delta\alpha) = 3-3\delta\alpha$. Thus a sufficient condition is, $3-3\delta\alpha \geq 2-3\delta+2\delta\alpha \Leftrightarrow 1 \geq 2\delta\alpha-3\delta(1-\alpha)$, which always holds since $\delta(1-\alpha) \geq \frac{1}{3}$.

Case 2 $\rho = 2\tau, 1 \leq \tau \leq M-2$. Player 2 has the move. The expected value of her payoffs are:

$$V_2^\rho(d_\rho|\nu_1(\cdot|\alpha, \delta, \pi_1)) = \rho + 1, \quad V_2^\rho(r_\rho|\nu_1(\cdot|\alpha, \delta, \pi_1)) = \delta(1-\alpha)(M+1) + \delta\alpha\rho + (1-\delta)(M+1).$$

Thus r_ρ is a best response provided: $(1-\delta\alpha)(M+1) + \delta\alpha\rho \geq \rho + 1 \Leftrightarrow (1-\delta\alpha)(M-\rho) \geq \delta\alpha$.

Since $M-\rho \geq 2$, a sufficient condition for r_ρ to be a best response is $2-2\delta\alpha \geq \delta\alpha \Leftrightarrow \delta\alpha \leq \frac{2}{3}$. Now $\delta(1-\alpha) \geq \frac{1}{3} \Rightarrow (1-\delta) + \delta\alpha \leq \frac{2}{3}$. Thus r_ρ is a best response under the given assumptions.

Having considered all possible cases we have established the inductive step. Thus there exists an equilibrium in which cooperation continues until the final node when $\delta(1-\alpha) \geq \frac{1}{3}$. Moreover if $\delta(1-\alpha) > \frac{1}{3}$ then r_ρ is the only best response at the relevant nodes, which establishes uniqueness of the equilibrium. \blacksquare

Proof of Proposition 5.2 We shall proceed by (backward) induction. At node the final node M , d_M is a dominant strategy. Now consider the decision at node $M-1$. Assume this node is off the equilibrium path. (This assumption will be confirmed when the proof is complete.)

Recall that, at nodes off the equilibrium path, the GBU updated preferences may be represented by the function, $W(a) = (1-\alpha)\max u(a) + \alpha\min u(a)$. Player 1 moves at this node. Her (Choquet) expected payoffs from continuing are:

$$\begin{aligned} V_1^{M-1}(r_{M-1}|\nu_1(\cdot|\alpha, \delta, \pi_1)) &= (1-\alpha)(M+2) + \alpha(M-1) = M+2-3\alpha < M \\ &= V_1^{M-1}(d_{M-1}|\nu_1(\cdot|\alpha, \delta, \pi_1)), \quad \text{since } \alpha \geq \frac{2}{3} \text{ we may conclude that } d_{M-1} \text{ is a best response.} \end{aligned}$$

Inductive step The inductive hypothesis is that d_κ is a best response at all nodes $M - 1 > \kappa > \rho > 1$. Now consider the decision at node ρ . Assume this node is off the equilibrium path and that player i has the move at node ρ . His/her expected payoffs are given by:

$$V_i^\rho(d_\rho|\nu_1(\cdot|\alpha, \delta, \pi_i)) = \rho + 1, \quad V_i^\rho(r_\rho|\nu_1(\cdot|\alpha, \delta, \pi_i)) = (1 - \alpha)(\rho + 3) + \alpha\rho = \rho + 3(1 - \alpha).$$

(Player i perceives no ambiguity about his/her own move at node $\rho + 2$.) Thus d_ρ is a best response provided: $1 \geq 3(1 - \alpha) \Leftrightarrow \alpha \geq \frac{2}{3}$. This establishes by induction that d_ρ is a best response at all nodes ρ such that $1 < \rho \leq M$, provided they are off the equilibrium path.

Node 1 Finally we need to consider the initial node, which is different since it is on the equilibrium path. Player 1 has to move at this node. Her expected payoffs are:

$$V_1^1(d_1|\nu_1(\cdot|\alpha, \delta, \pi_1)) = 2, \quad V_1^1(r_1|\nu_1(\cdot|\alpha, \delta, \pi_1)) = \delta(1 - \alpha)4 + \delta\alpha + (1 - \delta) = 3\delta(1 - \alpha) + 1.$$

Since $\alpha \geq \frac{2}{3}$ implies $\delta(1 - \alpha) \leq \frac{1}{3}$, which implies that d_1 is a best response at the initial node. This confirms our hypothesis that subsequent nodes are off the equilibrium path. The result follows. ■

Proposition A.1 For $\alpha < \frac{2}{3}$ and $\delta(1 - \alpha) < \frac{1}{3}$, Γ does not have a singleton CP-EUA.

The proof of Proposition A.1 follows from Lemmas A.3, A.4 and A.5.

Lemma A.3 Assume $\alpha < \frac{2}{3}$ then at any node $\tau, M > \tau \geq 2$, which is off the equilibrium path, the player to move at node τ will choose to play right, i.e. r_τ .

Proof. We shall proceed by (backward) induction. To start the induction consider the decision at node $M - 1$.

Node $M - 1$ Assume this node is off the equilibrium path. Player 1 has the move. Her expected payoff from continuing is, $V^1(r_{2M-1}|\nu_1(\cdot|\alpha, \delta, \pi_1)) = (1 - \alpha)(M + 2) + \alpha(M - 1) = M + 2 - 3\alpha \geq M = V^1(d_{M-1}|\nu_1(\cdot|\alpha, \delta, \pi_1))$, since $\alpha < \frac{2}{3}$.

Inductive step Since τ is off the equilibrium path so are all nodes which succeed it. The inductive hypothesis is that r_κ is a best response at all nodes $M - 1 > \kappa > \rho \geq \tau$. Now consider the decision at node ρ . First assume that Player 2 has the move at node ρ , which implies that ρ is an even number. His expected payoffs are given by:

$$V^2(d_\rho|\nu_2(\cdot|\alpha, \delta, \pi_2)) = \rho + 1, \quad V^2(r_\rho|\nu_2(\cdot|\alpha, \delta, \pi_2)) = (1 - \alpha)(M + 1) + \alpha\rho.$$

Thus r_ρ is a best response provided: $(1 - \alpha)(M + 1 - \rho) \geq 1$. Note that $\rho \leq M - 2$ and since $\alpha < \frac{2}{3}$, $3(1 - \alpha) > 1$ hence $(1 - \alpha)(M + 1 - \rho) \geq 1$ which establishes that right is a best response in this case.

Now assume that Player 1 has the move at node ρ . Her expected payoffs are given by:

$$V_1^\rho(d_\rho | \nu_1(\cdot | \alpha, \delta, \pi_1)) = \rho + 1, \quad V_1^\rho(r_\rho | \nu_1(\cdot | \alpha, \delta, \pi_1)) = (1 - \alpha)(M + 2) + \alpha\rho.$$

The analysis for Player 2 shows that r_ρ is also a best response in this case.

This establishes the inductive step. The result follows. ■

Lemma A.4 *Assume $\alpha < \frac{2}{3}$ and $\delta(1 - \alpha) < \frac{1}{3}$. Let Γ be a M -stage centipede game, where $M \geq 4$. Then there does not exist a singleton equilibrium in which Player 1 plays d_1 at node 1.*

Proof. Suppose if possible that there exists a singleton equilibrium in which Player 1 plays d_1 at node 1. Then all subsequent nodes are off the equilibrium path. By Lemma A.3 players will choose right, r_ρ , at such nodes. Given this Player 1's expected payoffs at node 1 are:

$$V_1^1(d_1 | \nu_1(\cdot | \alpha, \delta, \pi_1)) = 2,$$

$$V_1^1(r_1 | \nu_1(\cdot | \alpha, \delta, \pi_1)) = \delta(1 - \alpha)(M + 2) + \delta(1 - \alpha) \cdot 1 + (1 - \delta)(M - 1).$$

Thus d_1 is a best response if $2 \geq (1 - \delta\alpha)M + 3\delta(1 - \alpha) + \delta - 1$

$$\Leftrightarrow 3 - 3\delta(1 - \alpha) - \delta \geq (1 - \delta\alpha)M \geq 4 - 4\delta\alpha, \quad \text{since } M \geq 4.$$

$\Leftrightarrow 7\delta\alpha - 4\delta \geq 1$. Since $\alpha < \frac{2}{3}$, $1 > \frac{2}{3}\delta = \frac{14}{3}\delta - 4\delta > 7\delta\alpha - 4\delta$, hence d_1 cannot be a best response at node 1. ■

Lemma A.5 *Assume $\alpha < \frac{2}{3}$ and $\delta(1 - \alpha) < \frac{1}{3}$. Let Γ be a M -stage centipede game, where $M \geq 4$. Then there does not exist a singleton equilibrium in which Player 1 plays r_1 at node 1.*

Proof. Suppose if possible such an equilibrium exists. Let τ denote the first node at which a player fails to cooperate. Since cooperation will definitely break down at or before node M , we know that $1 < \tau \leq M$. There are three possible cases to consider.

Case 1 $\tau = M$. Consider the decision of Player 1 at node $M - 1$. Her expected payoffs are:

$$V_1^{M-1}(d_{M-1} | \nu_1(\cdot | \alpha, \delta, \pi_1)) = M,$$

$$\begin{aligned} V_1^{M-1}(r_{M-1} | \nu_1(\cdot | \alpha, \delta, \pi_1)) &= \delta(1 - \alpha)(M + 2) + \delta\alpha(M - 1) + (1 - \delta)(M - 1) \\ &= M + 3\delta(1 - \alpha) - 1. \end{aligned}$$

However $\delta(1 - \alpha) < \frac{1}{3}$, implies that r_{M-1} is not a best response. Thus it is not possible that $\tau = M - 1$.

Case 2 $\tau \leq M - 1$ and $\tau = 2k + 1$ This implies that Player 1 moves at node τ and that all subsequent nodes are off the equilibrium path. By Lemma A.3, right is a best response at these nodes. Hence her expected payoffs are:

$$V_1^\tau(d_\tau | \nu_1(\cdot | \alpha, \delta, \pi_1)) = \tau + 1, \quad V_1^\tau(r_\tau | \nu_1(\cdot | \alpha, \delta, \pi_1)) = \delta(1 - \alpha)(M + 2) + \delta\alpha\tau + (1 - \delta)(M - 1).$$

Thus d_τ is a best response provided,

$$\tau + 1 \geq (1 - \delta\alpha)M + 2\delta(1 - \alpha) + \delta\alpha\tau - (1 - \delta) \Leftrightarrow 2 + 2\delta\alpha - 3\delta \geq (1 - \delta\alpha)(M - \tau).$$

Since $M - \tau \geq 3$, a necessary condition for d_τ to be a best response is:

$5\delta\alpha - 3\delta \geq 1 \Leftrightarrow 2\delta\alpha \geq 1 + 3\delta(1 - \alpha)$. However the latter inequality cannot be satisfied since $\alpha < \frac{2}{3}$, $3\delta(1 - \alpha) > \delta \geq \delta\alpha$, and $\delta\alpha < 1$. Hence it is not possible that τ is odd.

Case 3 $\tau \leq M - 2$ and $\tau = 2k$ This implies that Player 2 moves at node τ and that all subsequent nodes are off the equilibrium path. His payoffs are

$$V_2^\tau(d_\tau | \nu_2(\cdot | \alpha, \delta, \pi_2)) = \tau + 1, \quad V_2^\tau(r_\tau | \nu_2(\cdot | \alpha, \delta, \pi_2)) = \delta(1 - \alpha)(M + 1) + \delta\alpha\tau + (1 - \delta)(M + 1).$$

For d_τ to be a best response we need

$$\tau + 1 \geq (1 - \delta\alpha)M + 1 - \delta\alpha + \delta\alpha\tau \Leftrightarrow \delta\alpha \geq (1 - \delta\alpha)(M - \tau).$$

Since $\alpha < \frac{2}{3}$, $\delta\alpha < \frac{2}{3}$ and $(1 - \delta\alpha) > \frac{1}{3}$. In addition $M - \tau \geq 2$, which implies that this inequality can never be satisfied. Thus τ cannot be even.

Having considered all possibilities we may conclude that there is no pure equilibrium where Player 1 plays r_1 at node 1. ■

Proposition 5.3 *If $\alpha < \frac{2}{3}$ and $\delta(1 - \alpha) < \frac{1}{3}$, there is a CP-EUA in which the support of Player 1's (resp. 2's) beliefs is $\{s_2^M, s_2^{M-2}\}$ (resp. $\{s_1^{M-1}, s_1^{M+1}\}$). The game will end at $M - 2$ or M with Player 2 exiting or at $M - 1$ with Player 1 exiting.*

Proof. Assume that Player 1's beliefs are a neo-additive capacity based on the additive probability π_1 defined by $\pi_1(s_2^M) = p$, for $\pi_1(s_2^{M-2}) = 1 - p$, $\pi_1(s_2) = 0$ otherwise, where $p = \frac{\delta(2-3\alpha)}{(1-\delta)}$. Likewise assume that Player 2's beliefs are a neo-additive capacity based on the additive probability π_2 defined by $\pi_2(s_1^{M+1}) = q$, for $\pi_2(s_1^{M-1}) = 1 - q$, $\pi_2(s_1) = 0$ otherwise, where $q = \frac{1-3\delta(1-\alpha)}{3(1-\delta)}$.

First consider the updated beliefs. At any node ρ , $0 \leq \rho \leq M - 2$, the updated beliefs are a neo-additive capacity with the same δ and α . The new probability π' is the restriction of the prior probability π to the set of continuation strategies. At node $M - 1$, the GBU rule implies the updated beliefs are a neo-additive capacity with the same α , the updated δ given by $\delta'_1 := \frac{\delta}{\delta + (1-\delta)p}$. The

updated π assigns probability 1 to s_2^M .²⁸ We do not need to specify the beliefs at node M . Player 2 has a dominant strategy at this node which he always will choose.

Player 2 Consider Player 2's decision at node $M - 2$. His actions yield payoffs,

$V_2^{M-2}(d_{M-2}|\nu_2(\cdot|\alpha, \delta, \pi_2))$ and $V_2^{M-2}(r_{M-2}|\nu_2(\cdot|\alpha, \delta, \pi_2))$. For both to be best responses they must yield the same expected payoff, which occurs when $q = \frac{1-3\delta(1-\alpha)}{3(1-\delta)}$, by equation (2). Since $\delta(1-\alpha) \leq \frac{1}{3}$ by hypothesis, $q \geq 0$. In addition $q \leq 1 \Leftrightarrow 3(1-\delta) \geq 1-3\delta(1-\alpha)$, which holds since $\alpha < \frac{2}{3}$ implies $(1-\alpha) > \frac{1}{3}$ thus $1-3\delta(1-\alpha) < 3-3\delta$.

Player 1 Now consider Player 1's decision at node $M - 1$. His updated δ is given by: $\delta'_1 = \delta(p | m') := \frac{\delta}{\delta+(1-\delta)p}$. His strategies yield payoffs $V_1^{M-1}(d_{M-1}|\nu_1(\cdot|\alpha, \delta, \pi_1))$ and $V_1^{M-1}(r_{M-1}|\nu_1(\cdot|\alpha, \delta, \pi_1))$. For both of these strategies to be best responses they must yield the same expected utility which occurs when $p = \frac{2\delta-3\delta\alpha}{1-\delta}$, by equation (3). Since $\frac{2}{3} > \alpha$ by hypothesis, $p \geq 0$. Moreover $1 \geq p \Leftrightarrow 1-\delta \geq 2\delta-3\delta\alpha \Leftrightarrow \frac{1}{3} \geq \delta(1-\alpha)$, which also holds by hypothesis. ■

A.3 Bargaining

Proof of Proposition 6.1 To establish that the strategy profile specified in the statement of Proposition 6.1 constitutes a CP-EUA, we work backwards from the end of the game.

At $t = 3$. Since this is the last period, it is a best response for player 2 to accept *any* proposed division $\langle z, 1-z \rangle$ by player 1. So in any consistent planning equilibrium, player 1 will propose a $\langle 1, 0 \rangle$ division after any history.

At $t = 2$. If player 1 rejects player 2's proposed division $\langle y, 1-y \rangle$, then the game continues to period 3 where we have established player 1 gets the entire cake. However, the neo-expected payoff to player 1 in period 2 of *rejecting* 2's offer is $(1-\alpha)\beta$. This is because by rejecting player 2's proposed division, player 1 is now in an "off-equilibrium" history h , with $\delta(h) = 1$. Thus her neo-expected payoff is simply an $(\alpha, 1-\alpha)$ -weighting of the best and worst outcome that can occur (given his own continuation strategy). Hence for player 2's proposed division $\langle y, 1-y \rangle$ to be accepted by player 1 requires $y \geq (1-\alpha)\beta$.

So the neo-expected payoff for player 2 of proposing the division $\langle (1-\alpha)\beta, 1-(1-\alpha)\beta \rangle$ in round 2 is $(1-\delta_2)(1-(1-\alpha)\beta) + \delta_2(1-\alpha)\max\{1-(1-\alpha)\beta, \beta\}$, where $\delta_2 = \delta/(\delta+(1-\delta)\pi)$ if $\langle x, 1-x \rangle = \langle 1, 0 \rangle$, =1 otherwise. Notice that the 'best' outcome for player 2 is either $(1-(1-\alpha)\beta)$

²⁸The form of the updated beliefs follows from the formulae in Eichberger, Grant, and Kelsey (2010).

or possibly β which is what she would secure if player 1 rejected his proposed division and then followed that by proposing in period 3 the (extraordinarily generous) division $\langle 0, 1 \rangle$.

If player 2 proposes $\langle y, 1 - y \rangle$ with $y < (1 - \alpha)\beta$, then according to our putative equilibrium, player 1 rejects and player 2 receives a payoff of zero. However, his neo-expected payoff of that is $(1 - \alpha) \max\{1 - y, \beta\}$, where $1 - y$ is what he would get if player 1 actually accepted his proposed division. Thus his best deviation is the proposal $\langle y, 1 - y \rangle = \langle 0, 1 \rangle$.

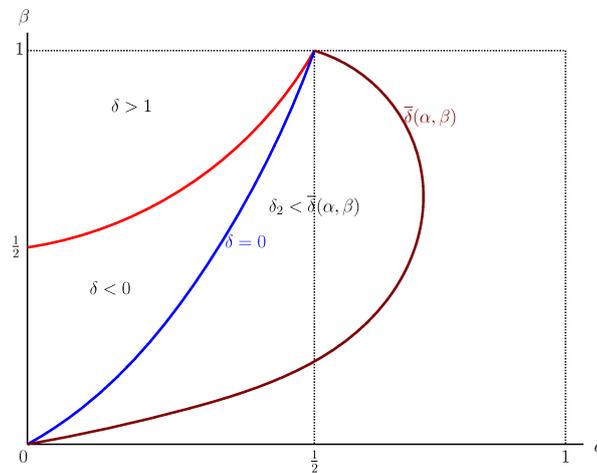
So we require:

$$(1 - \delta_2)(1 - (1 - \alpha)\beta) + \delta_2(1 - \alpha) \max\{1 - (1 - \alpha)\beta, \beta\} \geq (1 - \alpha). \quad (6)$$

This is equivalent to

$$\bar{\delta}(\alpha, \beta) := \frac{\alpha - (1 - \alpha)\beta}{1 - (1 - \alpha)[\max\{1 - (1 - \alpha)\beta, \beta\} + \beta]} \geq \delta_2.$$

The following figure shows the region of parameters (α, β) satisfying this constraint.



parameter region: $\delta_2 \leq \bar{\delta}(\alpha, \beta)$

Clearly, when $\delta_2 = 1$ (that is, when we are already off the equilibrium path), (6) cannot hold for any α strictly less than 1. So in that case, player 2's best action is indeed to propose $\langle y, 1 - y \rangle = \langle 0, 1 \rangle$, which is what our putative equilibrium strategy specifies in these circumstances.

At $t = 1$, if player 2 rejects player 1's proposed division $\langle x, 1 - x \rangle$, then the game continues to $t = 2$ in which case he offers player 1 the division $\langle y^*, 1 - y^* \rangle = \langle (1 - \alpha)\beta, 1 - (1 - \alpha)\beta \rangle$, which according to her strategy she accepts.

On the equilibrium path, $\delta_2 = \delta < 1$, the neo-expected payoff for player 2 in period 1 of rejecting player 1's proposed division $\langle x, 1 - x \rangle$ is $(1 - \delta)(1 - (1 - \alpha)\beta)\beta + \delta\alpha$. In this putative equilibrium Player 1 proposes the division $\langle x^*, 1 - x^* \rangle = \langle 1, 0 \rangle$ which player 2 rejects, since she anticipates an offer of $\langle (1 - \alpha)\beta, 1 - (1 - \alpha)\beta \rangle$ in period 2. So her period 1 neo-expected payoff is $(1 - \delta)(1 - \alpha)\beta^2 + \delta(1 - \alpha)$, since the 'best' thing that could happen is player 2 actually accepts her proposed division $\langle 1, 0 \rangle$ today!

Alternatively, player 1 might propose a division $\langle \tilde{x}, 1 - \tilde{x} \rangle$, that would just entice player 2 to accept today. That is, she could propose $\langle \tilde{x}, 1 - \tilde{x} \rangle$ where

$$\tilde{x} = 1 - [(1 - \delta)(1 - (1 - \alpha)\beta) + \delta(1 - \alpha)\beta \max\{1 - (1 - \alpha)\beta, \beta\}].$$

The neo-expected payoff for player 1 of this deviation is: $\delta(1 - \alpha)\max\{\tilde{x}, \beta\}$. In order for it not to be "profitable", we require $(1 - \delta)(1 - \alpha)\beta^2 + \delta(1 - \alpha) \geq \delta(1 - \alpha)\max\{\tilde{x}, \beta\}$, which holds since $\max\{\tilde{x}, \beta\} < 1$. ■