

Cognitive hierarchies and the centipede game*

Margherita Bottero[†]

SSE/EFI Working Paper Series in Economics and Finance, No 723

27 August 2010

Abstract

In the present work, I adopt the cognitive hierarchy approach to analyze the centipede game. To this end, I present and study an extensive-form version of Camerer et al.'s (2004) original normal-form model. The resulting predictions are evaluated empirically using laboratory data borrowed from a previously published experiment. The paper features two main contributions. First, it presents a parsimonious model that can, in principle, be generalized to any two-person extensive-form game of perfect information. Secondly, it demonstrates that in the centipede game the cognitive hierarchy approach leads to predictions which are not fully backwardly inductive and that can help to explain some key feature of the experimental data.

Keywords: centipede game, cognitive hierarchy, paradox backward induction, experimental data analysis

JEL classification: C72, C92, D03, D81

*I would like to thank Tore Ellingsen, Erik Mohlin, Juanna Joensen, Björn Wallace, Robert Östling, Jörgen Weibull and participants at the AISC VI meeting in Naples, at the SSE lunch seminar series and at the 3rd Maastricht Behavioral and Experimental Economics Symposium for their helpful comments and suggestions. I gratefully acknowledge the Knut and Alice Wallenberg Research Foundation for financial support.

[†]E-mail: margherita.bottero@hhs.se. Mail: Department of Economics, Stockholm School of Economics, P.O. Box 6501, SE-113 83 Stockholm, Sweden. Tel.: +46 (0) 762353874. Fax: +46 (8) 313207.

1 Introduction

There is by now ample evidence demonstrating that when individuals are presented with game theoretic interactions they frequently opt for behaviors different from the equilibrium predictions, especially during the first rounds of play (see for instance Costa-Gomes and Crawford, 2006, and references therein). This empirical anomaly motivates the ongoing search for new, non-equilibrium, models, with the goal of explaining the behavior of actual players as observed in laboratory or field experiments. Among the many contributions in this literature, the Level k model (Nagel, 1995; Stahl and Wilson, 1994, 1995) and its generalization, the cognitive hierarchy model (Camerer et al., 2004), stand out for their success in predicting and explaining actual behavior in a number of normal-form games, relying only on a parsimonious set of assumptions¹. How these models perform when applied to extensive-form games is, however, still an open question. Here I address this issue by suggesting and testing a possible extension of the cognitive hierarchy model that is directly applicable to two-player extensive-form games of perfect information. Within this class of games, I focus on the well-known centipede game (Rosenthal, 1982).

The centipede game represents an interactive situation involving two players that play sequentially and non-cooperatively for a limited number of rounds. The players take turns in deciding whether to terminate the game or to continue it, thereby passing this decision on to the opponent. Continuing increases the value of the joint payoffs, the pot, but it is not individually profitable if the game is terminated immediately thereafter. If that is the case, the player who passed receives a payoff lower than what he could have secured himself by terminating at the preceding round. If players seek to maximize their monetary gains, the centipede game features a unique (subgame perfect) equilibrium, in which the player who moves first terminates the game. Such a sharp theoretical prediction is at odds with the behavior of actual players, who typically continue for some rounds and occasionally until the end. No, or very few, participants have been observed taking the pot at the first round (see for instance Fey et al., 1996; McKelvey and Palfrey, 1992; Nagel and Tang, 1998; Rapoport et al., 2003).

This divergence between the equilibrium predictions and the experimental evidence appears analogous to that successfully addressed by the non-equilibrium models in normal-form games. For instance, the Level k model was first envisioned to explain why players' behavior in the n -person guessing game differed substantially from the equilibrium prescriptions (Nagel, 1993). This seminal contribution was soon generalized, extended and profitably applied to a large number of other games (Stahl and Wilson 1994, 1995; Costa-

¹See for instance Costa-Gomes et al. (2009) for various normal-form game applications of the Level k and cognitive hierarchy model based on laboratory data. Crawford and Iriberri (2007a) and Brocas et al. (2009) study games with private information using the Level k model. Östling et al. (2009) and Brown et al. (2009) test the Level k models using field data. Ellingsen and Östling, (2006) have studied several games with pre-play communication, where the communication phase and the actual play of the game can be thought of as two separate rounds of an extensive-form game.

Gomez et al., 2001; Costa-Gomez and Crawford, 2006; Ho et al., 1998; Nagel, 1995).

The main novelty of the Level k model is to assume that players' behavior can be described by their (behavioral) type which they draw at the beginning of the game from a hierarchy of possible types. At the bottom of this hierarchy lies a non-strategic type 0, who uniformly randomizes among his actions. Next, there is type 1, who believes that the opponent will act as a type 0, and plays a best response to this conjecture. Higher types adjust their beliefs via recursive thought experiments: type 2 best responds to type 1, type 3 to type 2 etcetera. This assumption is advantageous because it grants rationality to players of all types but type 0, simultaneously allowing them to adopt non-equilibrium behaviors, which follow from the simplified belief structure described by the hierarchy.

Camerer et al.'s (2004) cognitive hierarchy model of games was presented as a generalization of the Level k approach. In particular, it allows type k players, $k > 0$, to have more complex beliefs on the type of the opponent by envisioning the possibility of him being any of the lower types, from $k-1$ to 0. Thus, a type k not only best responds to type $k-1$, but to a mixture of lower-level types. The types' frequencies may be either treated as a parametrized distribution, usually Poisson, or as an unconstrained type distribution. These richer assumptions, that nest the Level k , have been successfully employed to explain players' behavior in a number of two-person finite dominance-solvable games, such as guessing games (Costa-Gomez and Crawford, 2006), zero-sum hide-and-seek games (Crawford and Iriberri, 2007b) or even the normal-form centipede games (Kawagoe and Takizawa, 2010).

In this paper, I present a version of the cognitive hierarchy model that is directly applicable to extensive-form two-person games of perfect information, focussing on the six-step centipede protocol studied by McKelvey and Palfrey (1992). I call this model the *EF-CH* model, where the acronym stands for extensive-form cognitive hierarchy. Possibly, the *EF-CH* is the most faithful transposition of Camerer's model to extensive-form games. The choice to apply the model to McKelvey and Palfrey's six-step centipede game is functional to produce predictions that are readily testable, thanks to the experimental data collected and made available by the authors. The model, however, can easily be generalized to any finite extensive-form game.

The *EF-CH* model rests on four assumptions. First, players can be of one of three types, 0, 1 and 2. Second, type 0 players randomize between actions with a constant probability, while type 1 and 2 ones use backward induction. Third, players of type 1 and 2 think that the opponent is of any of the type below their own, and they share the same belief on type 0s' randomization probability. Note that type 1s and type 2s do not envision the possibility of types equal or higher than their own. The fourth assumption is methodological, players maximize their monetary payoffs. Together, these hypotheses allow me to derive the payoff-maximizing behavior, in expectation, for each type and thus make verifiable predictions on the play of the game and the frequency of types in the experimental population.

In addition, I also study two extensions of such a model. First, I allow type 1 and type 2 players to have different beliefs on the behavior of type 0s. This special case, which I call *heterogenous EF-CH* (*HEF-CH*), features a more "sophisticated" type 2, that not only forms his own belief on the behavior of type 0s, but also holds a (second order) belief on type 1s' beliefs on type 0. This novel assumption departs from the original spirit of the cognitive hierarchy model, undermining the tight recursive structure of the hierarchy, but it allows for a more flexible structure that can better accommodate a larger variety of situations. For instance, imagine that players in a population can be divided in three groups. Players in the first group randomize, while those in the second group expect, erroneously, to meet altruistic opponents and best respond to that belief. Players in the third group instead realize the possibility to meet players from both the previous groups, and correctly anticipate and best respond to their behavior. For a practical example, it is possible to think of situations involving different sets of social norms that come to coexist. This and similar situations can be captured by the *HEF-CH* specification. Applied to the previous example, the model would let type 1 players believe that type 0s are altruists, and allow type 2 players to correctly recognize type 1s' mistake.

The second extension I consider is to allow Bayesian updating in the players' decision process. There is evidence that when playing repeatedly in the same interaction players revise their prior beliefs on the opponent using Bayes' Law, although this topic is still open to debate (Cosmides and Tooby (1996), Fudenberg and Levine (1998) and references therein). Updating, in fact, is a cognitively advanced task, which requires the use of cognitive resources to improve upon the priors. For this reason, and to be consistent with the spirit of the cognitive hierarchy, I allow only type 2 players to update their belief on the type of the opponent.

The *EF-CH* model produces a large number of theoretical results. In particular, it shows, for any belief the players may have, the behaviors that maximize their expected payoff. Among these, there is the backward induction result that to terminate at the first round, but it is far from being the only one. Other beliefs are such it is payoff maximizing for type 1 and type 2 players to continue to the last node of the centipede. Bayesian updating does not have any observable effect on the predictions of the model, neither by introducing new payoff-maximizing behaviors nor eliminating existing ones. However, it alters the set of beliefs for which the prescribed behaviors are payoff-maximizing. Using the assumptions of the *HEF-CH* case, instead, I can characterize new payoff-maximizing behaviors, including type 2 players terminating at the central round. Using data from McKelvey and Palfrey's (1992) experiment, I test whether these predictions can actually explain the experimental behavior. I find that the *EF-CH* model predicts the experimental behavior to a good extent. However, the heterogenous extension outperforms the *EF-CH* model in terms of the empirical fit. In general, all the different theoretical predictions fall short of fully fitting the data, failing to predict enough terminations at the central nodes of the game. I attribute this shortcoming to the assumption of constant randomization

from part of type 0s, and discuss it further in section 6.

To my knowledge, this work is the first that applies and tests the cognitive hierarchy framework directly to the extensive-form centipede game. Camerer (2003) mentions that Level k could explain the experimental evidence relative to this game, but he does not provide a formal analysis. Ho et al. (2008) study a normal-form centipede game assuming a cognitive hierarchy structure that allows for learning. Similarly, Nagel and Tang (1998) look at a normal-form centipede game to investigate different learning models. Kawagoe and Takizawa (2010) compare the empirical fit of Level k model, the Agent quantal response model² (Fey et al., 1996; McKelvey and Palfrey, 1998) and the Agent quantal response model with altruistic types (McKelvey and Palfrey, 1992, 1998) when applied to the reduced normal-form³ of the centipede game. The results of their analysis show that the Level k model outperforms the other two models considered. Compared to that of Kawagoe and Takizawa, my work deals with the cognitive hierarchy model and aims at evaluating its reach when applied directly to the extensive-form of the centipede game. This approach has the side advantage of allowing me to consider directly the possibility of Bayesian updating.

Palacios-Huerta and Volij (2009) have recently presented interesting experimental results on the centipede game (for a discussion, see also List et al., 2010). The authors have documented that chess players rarely fail to terminate at the first round of the game (the subgame perfect equilibrium) when playing with other chess players, while they typically continue to later rounds when playing with college students. In line of principle, this result could be given a cognitive hierarchy interpretation. It is possible that chess players have beliefs on the behavior of other chess players and college students such that it is payoff maximizing, in expectation, to terminate as soon as possible with the first ones and to continue for a while with the latter ones⁴. Finally, in a recent paper (2010), Gerber and Wichardt investigate the efficacy of bonuses and insurance as deferring devices to postpone termination in the play of the centipede game⁵. In explaining the experimental results, the authors discuss a model that has a flavour of non-equilibrium thinking, however they are careful in stating that they are not concerned with investigating the reach of Level k in the centipede game. My work and that of Gerber and Wichardt can benefit from one another, however they address two different questions. I am here concerned with the issue of how to apply the cognitive hierarchy model to the centipede game.

The remainder of the paper is organized as follows. Section 2 discusses the centipede

²The agent quantal response model assumes that each player treats his future self as an independent player with a known probability distribution over actions.

³To corroborate this assumption, the authors present the result of an experiment showing that players do not behave differently when the same game is presented to them in extensive or normal form.

⁴Interestingly, also college students are documented to terminate earlier when playing with chess players rather than with other students.

⁵The authors are motivated by the fact that, in the standard centipede game, longer interactions result in increased aggregate welfare for both players.

game and the relevant literature in ampler detail. Section 3 introduces the assumptions of the model and section 4 analyzes the results. Section 5 presents the empirical analysis. Section 6 discusses the results and section 7 concludes.

2 The centipede game

In game theory, the centipede is a finite and sequential game of perfect information with two player roles, here labelled *A* and *B*. Following McKelvey and Palfrey (1992), in this work I consider the 6-node version of the centipede game that is displayed in Figure 1. At each node j , the player who has to move can choose either to continue to play (c_j) or to terminate the game (t_j). Upon the decision to terminate the game at a node j , the players receive the payoffs appearing at the end of branch t_j , where the first amount pertains to player (role) *A*. At the first node, the total pot sums up to \$0.5, thereafter, each time it is passed on its total value increases exponentially by a factor of two, until it reaches a total of \$32 in the last node.




Fig. 1. The Centipede game

Since the game is sequential and there is perfect information about the history of play, a new subgame begins at each node, for a total of six (proper) subgames. Subgames are important in the search for the equilibrium because a strategy profile such that the players do not want to deviate at any subgame is also a subgame perfect strategy profile (Selten, 1965). As it is standard in game theory, it is assumed that the players are rational and have Von Neumann Morgenstern utility functions. Kuhn (1950, 1953) demonstrated that subgame perfect equilibria can be found working *backwardly* from the end of the tree, replacing the starting node of each subgame by the vector of payoffs that corresponds to the optimal choice at the subsequent nodes.

Applied to the game in Figure 1, this procedure begins with considering player *B*'s choice at node 6, the last subgame of the game. There, player *B*'s optimal choice is t_6 , because, being rational, he chooses the action delivering him the higher payoff. Then, the players can replace the subgame with the payoff vector corresponding to the choice t_6 , obtaining a new, artificial centipede. Applying the same reasoning inductively, it is possible to establish that player *A*, being rational, will choose t_5 at node 5, and so on. The

collection of these choices, one for each personal decision node of the game tree, defines the unique⁶ subgame perfect equilibrium of the game: to terminate as soon as the player gets the move.

Despite being straightforward, the backward induction result is uncomfortable, for it clashes with several laboratory observations and, admittedly, with the results of introspection. To improve on this front, the game theoretic literature has proposed several equilibrium refinements. In particular, the suggested models deal with the issue of what should players believe after observing a deviation from the equilibrium strategy. The backward induction procedure, in fact, rests on the assumption that both players know that both will keep the pot if given it, rather than pass it on, and both know that both know this fact, and so on. Modifying this assumption may allow to obtain other equilibria (see, for instance, Asheim and Dufwenberg, 2003; Ben-Porath, 1997; Kreps and Wilson, 1982; Kreps et al., 1982; Reny, 1998; Selten, 1975). In general, the refinements support the backward induction outcome, although allowing for different outcomes to be realized as well.

In this paper, I depart from equilibrium analysis, and only assume that players are rational, and that they hold mutual and subjective conjectures on the play of the opponent. To describe how these payoff-maximizing players play the game it is sufficient to know which conjecture they have about the opponent (a conjecture is here the node where the opponent is believed to terminate). In the next section, I will show that these simple requirements go a long way to explaining the laboratory evidence on the centipede.

As mentioned in the Introduction, the reason for choosing to focus on the game in Figure 1 is that it is part of a series⁷ of experiments conducted by McKelvey and Palfrey (1992). Thanks to the availability of their laboratory data, it will be possible to directly test the model's predictions. The experiment consists of three⁸ sessions, labeled 5, 6 and 7, two with twenty participants each, and one with eighteen, none of whom had played the game before. At the beginning of each session, participants were divided in two groups of ten (or nine in the case of session 6), and randomly allocated to role *A* or *B*. Each subject was then matched with an individual from the other group to play the game, for a total of ten games, or nine in the eighteen-player session. It was made common knowledge at the beginning of each session that no individual was going to be matched with the same opponent twice. In this sense, there is no learning, and matching can be treated as random.

The laboratory data that was collected in this experiment uncovered two main features that characterize the subjects' behavior, which I will discuss shortly. Table 1 shows the

⁶Thanks to the following proposition: every finite, generic, extensive form game of perfect information has a unique subgame perfect equilibrium.

⁷In the same paper, the authors consider also two four-step games, one with a "high payoff" treatment.

⁸Session 1, 2 and 3 involved other types of treatment. Specifically, these consisted of a shorter (4-step) centipede game, and a "high payoff" condition, with stakes substantially higher than those considered here.

frequency, f_j , of plays terminating at node j . The backward induction outcome, which would require $f_1 = 1$, never occurs in the data. What instead appears is that terminations occur mostly after the third node, and are concentrated around nodes 3 to 5. Few plays continue without interruption to the end of the centipede.

Table 1: Frequency of Observations at each Terminal Node

Session	N	f_1	f_2	f_3	f_4	f_5	f_6	f_7
5	100	.02	.09	.39	.28	.20	.01	.01
6	81	.00	.02	.04	.46	.35	.11	.02
7	100	.00	.07	.14	.43	.23	.12	.01
Total 5-7	281	.007	.064	.199	.384	.253	.078	.014

It is possible to use the f_j s to back out the implied probability p_j of terminating at any given node j , conditional to the event that node j is reached, $p_j = f_j / (f_j + \dots + f_7)$. These are shown in Table 2.

Table 2: Implied Take Probabilities

Session	p_1	p_2	p_3	p_4	p_5	p_6
5	0.02	0.09	0.44	0.56	0.91	0.50
6	0.00	0.02	0.04	0.49	0.72	0.82
7	0.00	0.07	0.15	0.54	0.64	0.92
Total 5-7	0.01	0.06	0.21	0.53	0.73	0.85

For instance, in Session 5, if a play reaches node 5, the ex ante probability it will end there is $p_5 = f_5 / (f_5 + f_6 + f_7) = 0.2 / (0.2 + 0.01 + 0.01) = 0.91$. The backward induction outcome would imply $p_1 = 1$, but it is never documented in the data. The most evident and consistent empirical pattern is that the probability of taking increases as the play approaches the last nodes⁹.

Summing up, there are two important empirical features that a theory aiming at explaining the centipede game should replicate. The first is the distribution of terminations concentrated around the mean, and the second is the pattern of the conditional take probability, which is increasing in the length of the centipede. The following three sections present, analyze and empirically evaluate the extent to which the cognitive hierarchy approach succeeds in explaining the centipede game.

⁹In Session 5, the figure at the last move (0.50) is based on two observations only.

3 The *EF-CH* model

As suggested by the analysis of players' behavior in several experiments, I assume that players can be one of three types, 0, 1 and 2, which are recursively defined as follows. Players of type 0 are at the bottom of the types' hierarchy, and behave as non-strategic randomizers. At the beginning of the game, they observe a realization q of the random variable \tilde{q} , and continue at each node with that probability. Let \tilde{q} be distributed with density f over the support $[0, 1]$. Note that there is no heterogeneity among type 0s, in the sense that they all randomize in the same way. Also note that when q is equal to $\frac{1}{2}$ we have the uniform randomization case. The next type defined in the hierarchy is type 1. Players of this type are Savage rational decision-makers that maximize their expected monetary payoff through reasoning by backward induction. In doing so, they assume that *all* the other players are of type 0, and they assign a point belief λ to q . The belief λ is the realization of a random variable $\tilde{\lambda}$ defined on the same support as q , which is observed by all type 1 players at the beginning of the game. This implies that, as was the case for type 0s, all players of type 1 are homogeneous in their beliefs (and hence in their behavior). The players of type 2 are at the top of the type hierarchy. They are similar to type 1s in three respects. First, they are Savage rational maximizers that reason by backward induction. Secondly, they observe the same realization λ and use it as point belief on q . Thirdly, they do not consider the possibility that other players may be of the same type as their own. However, differently from type 1s, they can guess correctly that a fraction s_0 of the opponents is of type 0 and a fraction s_1 is of type 1. Since the two fractions do not sum up to one¹⁰, type 2s normalize to $\sigma_0 = \frac{s_0}{s_0 + s_1}$ and to $\sigma_1 = 1 - \sigma_0 = \frac{s_1}{s_0 + s_1}$ the probability to be playing with a type 0 and a type 1 respectively.

I will extend this model in two ways. First, I will assume that type 2 players observe λ , and thus are able to correctly predict the behavior of type 1 players, but assign to q a point belief μ , which is the realization of a random variable $\tilde{\mu}$ different from $\tilde{\lambda}$ ¹¹. I call this case *HEF-CH*, which stands for *heterogenous EF-CH*, because type 1 and type 2 players have heterogeneous beliefs. The second extension that I consider is to allow for Bayesian updating during the play of the centipede game. In particular, I study an extension of the model above where I allow type 2 players to engage in Bayesian updating on the prior σ_0 , after having observed the opponent's move. Given that type 2s know the node at which type 1 players terminate, they update on σ_0 making use of the fact that *before that node* only players of type 0 drop off the game. Upon observing a continuation decision before the type 1s' expected terminal node, type 2 players re-assess downward the probability of being playing with a type 0. After observing the first pass from the opponent, the type 2 player can revise the prior to $\sigma_0^{R1} = \frac{\sigma_0 \rho_2}{1 - \sigma_0(1 - \rho_2)}$, where the

¹⁰I follow here Camerer et al. (2004), despite type 2 players know the true proportions of $C0$ and $C1$ in the population, they normalize them dividing by the sum, since $C2$ s do not envision the possibility of players of their own or higher type.

¹¹But defined on the same support $[0, 1]$.

superscript $R1$ stands for first revision. The updated probability of being playing with a type 1 becomes $\sigma_1^{R1} = 1 - \sigma_0^{R1} = \frac{1-\sigma_0}{1-\sigma_0(1-\rho_2)}$. If the opponent chooses to pass a second time, such a belief can be updated again, this time to $\sigma_0^{R2} = \frac{\sigma_0\rho_2^2}{1-\sigma_0(1-\rho_2^2)}$.

Given the length of the centipede game in Figure 1, players revise only twice (at nodes 3 and 5 for player role A and nodes 2 and 4 otherwise¹²). However, the updating process can be generalized to any number n of revisions: $\sigma_0^{Rn} = \frac{\sigma_0\rho_2^n}{1-\sigma_0(1-\rho_2^n)}$. As expected, the revised belief goes to 0 as the number of revisions goes to infinity

$$\lim_{n \rightarrow \infty} \sigma_0^{Rn} = \frac{\sigma_0\rho_2^n}{1-\sigma_0(1-\rho_2^n)} = 0$$

Similarly, σ_1^{Rn} goes to 1, $\lim_{n \rightarrow \infty} \sigma_1^{Rn} = \frac{1-\sigma_0}{1-\sigma_0(1-\rho_2^n)} = 1$.

The assumptions just described can be applied to any two-person extensive-form finite game of perfect information. The model will result, in general, in a series of *outcomes*, where by outcome I mean a collection of terminal nodes and beliefs for type 1 and 2, such that, given the beliefs, it is payoff maximizing to stop at the prescribed nodes. Note that *ex ante* players' beliefs can fall anywhere on the support, so that the model is descriptive of a variety of behaviors. To increase the precision of its predictions, it is possible to specify more accurate hypotheses on the players' beliefs. Alternatively, if one wants, or has to, remain agnostic about the players' beliefs, an "empirical approach" can be adopted. This consists in taking all the outcomes to the data and judge empirically which one fits the evidence better. In this paper, I will follow the latter approach. First, however, the next section presents the outcomes of the model.

4 Analyzing the model

In this section, I begin by studying the *EF-CH*'s predictions, then I consider in turn the *Bayesian* case, in which type 2 players undertake Bayesian updating, and the *HEF-CH* case, which allows type 2 players to form a belief μ on q that is different from λ . In general, the assumptions lead to a partition of the belief-space Ω in a number of non-overlapping components o such that $\bigcup_{o \in \Omega} o = \Omega$. Each component o , which I call *outcome*, is associated with a set of strategies (r_A^*, r_B^*) for type 1s and (t_A^*, t_B^*) for type 2s, such that if a player's belief(s) fall in that outcome it is payoff maximizing for him to follow r^* or t^* . For each outcome, or, alternatively, for each belief the players may have, it is possible to derive the probability \bar{q}_n , $n = 1, 2, \dots, 7$, that a randomly drawn pair of players terminates at node n . The probability \bar{q}_n is a function of the shares of the players types in the population, s_0 , s_1 and s_2 , and of q , the type 0's randomization probability, and, as we will see in the next section, will serve as the basis for the maximum likelihood analysis.

¹²Revision at node 6 is possible but useless, given that the game terminates there.

Let us begin with the *EF-CH* model, and assume that type 1s and type 2s have observed a realization λ of the belief on q . Consider a type 1 individual in role A that maximizes the expected payoff $E[\pi(\cdot)]$ of his actions. At the last node he compares the expected value of c_5 , $E\pi(c_5) = \lambda\pi(c_6) + (1 - \lambda)\pi(t_6)$, and that of t_5 , $\pi(t_5)$, and plays the action with highest expected value. Then, he takes one step back and proceeds the same way at node 3 and 1. Reasoning is a similar way for the other player role, the expected payoff maximizing strategies for a type 1 are

$$\begin{aligned} r_A^* &= \begin{cases} c_1c_3c_5 & \text{if } \lambda \geq \frac{1}{7} \\ t_1t_3t_5 & \text{otherwise} \end{cases} \\ r_B^* &= \begin{cases} c_2c_4t_6 & \text{if } \lambda \geq \frac{1}{7} \\ t_2t_4t_6 & \text{otherwise} \end{cases} \end{aligned} \quad (1)$$

By assumption, type 2 players assign to q the same λ as type 1, hence they can anticipate the behavior in equation 1. Assigning belief σ_0 to the event that the opponent is of type 0, type 2 players reason backwardly to find their payoff maximizing behavior. After some computations, it turns out that the payoff maximizing strategies t_A^*, t_B^* are

$$\begin{aligned} t_A^* &= \begin{cases} c_1c_3c_5 & \text{if } \sigma_0 \geq \bar{\lambda} \text{ and } \lambda \geq \frac{1}{7} \\ c_1c_3t_5 & \text{if } \sigma_0 < \bar{\lambda} \text{ and } \lambda \geq \frac{1}{7} \\ t_1t_3t_5 & \text{otherwise} \end{cases} \\ t_B^* &= \begin{cases} c_2c_4t_6 & \text{if } \lambda \geq \frac{1}{7} \\ t_2t_4t_6 & \text{otherwise} \end{cases} \end{aligned} \quad (2)$$

where $\bar{\lambda} = \frac{6.4}{22.4\lambda + 3.2}$. Interpreting the prescription in equation 2 is intuitive. Type 2s' strategy depends both on λ , which is the parameter describing the behavior of type 0s and type 1s, and on the belief σ_0 , on the shares of type 0s in the population. As long as the type 0 players are believed to randomize with a probability that is high, $\lambda \geq \frac{1}{7}$, it is optimal for the type 2s to terminate at the last round. In particular, if type 2s believe that there are many type 0s that continue with a high probability, then it is payoff-maximizing, in expectation, to pass even at node 5. Note however that the rationality assumption makes sure that, when playing in position B , type 2s always take at node 6. When $\lambda < \frac{1}{7}$, instead, both type 1s and type 2s terminate at the first node.

To sum up, the *EF-CH* assumptions can give rise to three outcomes, only one of which prescribes the backward induction behavior. The outcomes can be represented on

the (λ, σ_0) plane, as shown in Figure 2.




Fig. 2. Outcomes of the EF-CH model

Outcome i consists of those beliefs such that both types 1 and 2 continue to the end of the game. The beliefs that form outcome ii , instead, are such that it is payoff-maximizing for type 2 players to terminate at the last round, while for type 1s to continue to the end of the game. Finally, the beliefs in outcomes iii are such that both types stay out at the first round. Note that it is possible to read the behavior of type 1 players in correspondence of $\sigma_0 = 1$ ¹³. Figure 2 also portrays the uniform randomization hypothesis, $\lambda = \frac{1}{2}$. In this case, type 1 players play until the last round, while type 2 players may either continue to the last round or terminate at node 5 depending on their belief σ_0 . The termination probabilities are reported in Appendix B.

4.1 Introducing Bayesian updating

As mentioned in the Introduction and in section 3, one natural extension of the *EF-CH* model is to allow for Bayesian updating. Here I assume that type 2 players update their prior σ_0 on the share of type 0 players in the population. It turns out that if type 2 players make use of the updated beliefs σ_0^{R1} , σ_0^{R2} during the play of the game the outcomes of the Bayesian model are behaviorally indistinguishable from the non-Bayesian ones. To see why, consider first the case $\lambda < \frac{1}{7}$. Regardless of the belief σ_0 , it is optimal for type 2s to terminate at the first round. In this case, in fact, type 1s terminate at the first round (see equation 1) and type 0s are also believe to terminate with a high probability (λ is low). If $\lambda \geq \frac{1}{7}$, instead, the type 2s' payoff maximizing strategies remain the same as in equation 2, with the difference that the set of belief for which they are optimal changes slightly. In particular, $\bar{\lambda}$ is replaced by $\tilde{\lambda} = \frac{1}{7\lambda^3 - \lambda^2 + 1}$. The Bayesian outcomes Bi , Bii

¹³In fact, the assumptions of the model are such that it is "as if" players of type 1 had a belief $\sigma_0 = 1$ regarding the population composition.

and *Biii* are represented on the (λ, σ_0) plane in Figure 3.




Fig. 3. Outcomes of the Bayesian EF-CH model

With respect to the non-Bayesian model we can note two changes. First, for player role *A*, continuing at the last round is optimal for a smaller set of beliefs. This happens because when node 5 is reached a player has to believe in a high σ_0 to continue, because c_5 is profitable only with a type 0 opponent. However, during the play of the game σ_0 he has been steadily revised downwardly. Thus, it is payoff maximizing to play c_5 only for high prior σ_0 . It is also worth noting that when Bayesian updating is possible, playing in role *B* gives a strategic advantage over playing in role *A*. Starting to play at node 2 means in fact that the player can revise his prior *before* having to play the first time, simply by observing the opponent's move at node 1. While this does not presently affect the outcomes, it indicates a strategic advantage of being playing in position *B*.

4.2 The *HEF-CH* case

In the *HEF-CH* case, I relax the assumption that type 2s and type 1s hold the same belief on q . In particular, I assume that type 2 players assign a belief μ and correctly guess type 1s' belief λ . Thanks to this last assumption, we can consider separately the cases when $\lambda \geq \frac{1}{7}$, when type 1 players continue to the last round, and $\lambda < \frac{1}{7}$, when type 1s terminate at the first round (see equation 1). I begin with the case $\lambda \geq \frac{1}{7}$, which I call

$HEF-CH^+$ case. Type 2s' payoff maximizing strategies are given by

$$t_A^* = \begin{cases} c_1 c_3 c_5 & \text{if } \sigma_0 > \bar{\mu} \text{ and } \mu \geq \frac{1}{7} \\ c_1 c_3 t_5 & \text{if } \sigma_0 < \bar{\mu} \text{ and } \mu \geq \frac{1}{7} \text{ or } \mu < \frac{1}{7} \text{ and } \sigma_0 \leq \frac{2}{3} \\ c_1 t_3 t_5 & \text{if } \frac{2}{3} < \sigma_0 \leq 0.93 \text{ and } \mu < \frac{1}{7} \\ t_1 t_3 t_5 & \text{otherwise} \end{cases}$$

$$t_B^* = \begin{cases} c_2 c_4 t_6 & \text{if } \mu \geq \frac{1}{7} \text{ or } \mu < \frac{1}{7} \text{ with } \sigma_0 \leq \frac{2}{3} \\ c_2 t_4 t_6 & \text{if } \frac{2}{3} < \sigma_0 \leq 0.93 \text{ or } \mu < \frac{1}{7} \\ t_2 t_4 t_6 & \text{otherwise} \end{cases}$$

where $\bar{\mu} = \frac{6.4}{22.4\mu + 3.2}$. Compared with the $EF-CH$ model, the $HEF-CH^+$ case features one main difference, that there are now beliefs such that it is optimal for type 2 players to terminate at the central round. The reason for the appearance of a new outcome III is the following: termination at the central round is payoff maximizing for type 2s when $\lambda \geq \frac{1}{7}$, $\mu < \frac{1}{7}$ and σ_0 is high but not very high. If and only if these conditions are met, it is in expectation payoff maximizing to terminate at nodes 3 and 4. Since by definition in the $EF-CH$ case $\lambda = \mu$, the conditions for outcome III to be payoff-maximizing are never satisfied. Graphically, the four outcomes of the $HEF-CH^+$ can be represented as in Figure 4.




Fig. 4. Outcomes of the $HEF-CH^+$ case

Recall that type 2s know that type 1s will continue to the last round. Outcome I is identical to outcome i in the $EF-CH$ case. Outcome II is behaviorally identical to outcome ii for type 2 players, but note that it is here optimal for a larger set of beliefs. For instance, even when type 2 have a low belief on q , $\mu < \frac{1}{7}$, they may continue to the last round, if they believe that there are many type 1s in the population. As the share of type 1s

decreased, with $\mu < \frac{1}{7}$, type 2 players terminate increasingly early, at the central round, outcome *III*, or at the first round, outcome *IV*.

The second case to consider, which I call *HEF-CH⁻*, arises when $\lambda < \frac{1}{7}$. Here type 1 players terminate at the first round, according to equation 1. Note that this means a strategic advantage of being allocated to role *B*. Type 2 players in that role will get to know, before taking the first decision, the opponent's type. If the game is still on at node 2, then they are playing with type 0. It follows that type 2s' payoff maximizing strategy for position *B*, t_B^* , is the same as r_B^* (see equation 1). However, when playing in role *A*, a type 2 player has no information on the type of his opponent. Thus, if he believes $\mu < \frac{1}{7}$ it is optimal for him to play t_1 , because even if there were only types 0 in the population, such a μ_0 would be too low to make it profitable to play c_2 . If instead $\mu \geq \frac{1}{7}$, it turns out that it is payoff maximizing to play

$$\begin{aligned} t_A^* &= \begin{cases} c_1 c_3 c_5 & \text{if } \sigma_0 \geq \tilde{\mu} \\ t_1 t_3 t_5 & \text{otherwise} \end{cases} \\ t_B^* &= \begin{cases} c_2 c_4 t_6 & \text{if } \mu \leq \frac{1}{7} \\ t_2 t_4 t_6 & \text{otherwise} \end{cases} \end{aligned} \quad (3)$$

where $\tilde{\mu} = \frac{0.2}{0.6\mu + 2.4\mu^2 + 22.4\mu}$. As for the *HEF-CH⁺* model, we can represent graphically the possible outcomes on the (μ, σ_0) plane, holding $\lambda \geq \frac{1}{7}$. The outcomes are shown in Figure 5.




Fig. 5. Outcomes of the *HEF-CH⁻* model

Outcome *I⁻* consists of those beliefs that make it optimal to terminate at the last round. Outcome *II⁻* collects those beliefs such that $t_A^* = t_1 t_3 t_5$ and $t_B^* = c_2 c_4 t_6$. Finally, in outcome *III⁻*, the payoff-maximizing strategies are $t_A^* = t_1 t_3 t_5$ and $t_B^* = t_2 t_4 t_6$. The termination probabilities can be found in Appendix B.

It is possible to allow for Bayesian updating also in the *HEF-CH* case. The computation are shown in Appendix A. This extension produces three outcomes, of which only one, *BIII*, is behaviorally different from those already studied. According to *BIII*, type 1 players continue to the last round (terminal nodes c_5 and t_6). Type 2 players, instead, terminate at the first node when in position *A* and at node 6 when in position *B*.

5 Taking the model to the data

In this section, I use the data collected by McKelvey and Palfrey in their experiment (1992), to test the predictions of the *EF-CH* model and its extensions. In particular, I will use data from the whole sample, 281 observations. These were obtained dividing the 58 participants into three groups, and further splitting each group in two halves, assigning them to role *A* or *B*. Then, each player in a certain role played the game with all the participants in his group that were allocated to the other position. To address potential concerns of learning during these ten repetitions, I consider also a smaller sample with only first response data.

Before taking the theoretical results to the data, there are two considerations to be made. First, there are instances in which the behavior prescribed by different outcomes is indistinguishable. For instance, outcomes *i* and *Bi* prescribe to continue to the end of the game. Outcomes *ii* and *Bii* prescribe that type 2 players terminate at the last round, and that type 1 players continue to the last round. Outcomes *iii* and *III*⁻ prescribe that type 1s and 2s terminate at the first round. The impossibility to distinguish different outcomes by looking at players' behavior implies that I can not always infer the players' beliefs. The second consideration is that, given that I did not make any assumption on the players' beliefs, *ex ante* all the outcomes are equally likely. For this reason, here I will try to discriminate between (distinguishable) outcomes by searching empirically for the one that predicts and describes the experimental evidence best. To this end, I first employ a maximum likelihood analysis to find the parameter estimates of each outcome¹⁴. Then, I compare the different outcomes in two ways. First, I look at their information criteria AIC and BIC to shed some light on which specification is preferable from an econometric point of view. Secondly, I use the maximum likelihood estimated to derive each outcomes' predicted frequencies of termination and the implied take probabilities to compare them with the actual data. Finally, I estimate the model assuming that type 0 players actually exist in the population. Unreasonably high estimated shares of type 0s may suggest that the model is not well specified. Moreover, the estimation of q permits to gather more information on the way players randomize in the game.

Suppose now that type 2 and type 1 players have observed beliefs λ (and μ in the *HEF-CH* case) and σ_0 according to the model's hypotheses. Type 0 players follow the

¹⁴In the estimation, I will not allow for errors in actions. This mostly because McKelvey and Palfrey (1992) do not find evidence for that when analyzing the same dataset.

deterministic behavior described in section 3 and I make two further assumptions that will facilitate the empirical analysis. First, I assume that type 2s' beliefs on the population composition are correct in expectation. Secondly, I assume that there is the same distribution of types for each player role. In other words, regardless of whether a player is drawn to play in role A or B , he is equally likely to be of a certain type. Then, each repetition of the game can be looked at as an experiment with seven mutually exclusive ways to terminate, t_1, \dots, t_7 , which have probabilities $\bar{q}_1, \dots, \bar{q}_7 = 1 - \sum_{j=1}^6 \bar{q}_j$. These were derived for each outcome in the previous section. Repeating the experiment n independent times, and assuming that the \bar{q}_j s remain the same from trial to trial, the probability of observing a random sample $\mathbf{t} = [T_1, \dots, T_7 = n - \sum_{j=1}^6 T_j]$ is distributed according to a multivariate binomial

$$f(\mathbf{t}) = \frac{n!}{t_1! \dots t_6! \left(n - \sum_{j=1}^6 t_j\right)!} \bar{q}_1^{t_1} \dots \bar{q}_6^{t_6} \left(1 - \sum_{j=1}^6 \bar{q}_j\right)^{n - \sum_{j=1}^6 t_j} \quad (4)$$

The density function (4) depends on the parameters s_0, s_1, q , which are the population shares of players of type 0 and 1 and the probability q with which type 0s continue at each node. For any outcome, the likelihood function for a sample of N independent observations is the product of N densities

$$L(s_0, s_1, q) = \prod_{i=1}^{N+1} f(\mathbf{t}^i; s_0, s_1, q) \quad (5)$$

and the log-likelihood $\lambda = \log [L(s_0, s_1, q)]$ becomes

$$\lambda = \sum_{i=1}^{N+1} \ln A + \sum_{i=1}^{N+1} t_1^i \ln \bar{q}_1 + \dots + \sum_{i=1}^{N+1} t_6^i \ln \bar{q}_6 + \sum_{i=1}^{N+1} \left(n - \sum_{j=1}^6 t_j\right) \ln \left(1 - \sum_{j=1}^6 \bar{q}_j\right) \quad (6)$$

where $A := \frac{n!}{t_1! \dots t_6! \left(n - \sum_{j=1}^6 t_j\right)!}$.

By substituting the \bar{q}_j s implied by each outcome into expression (6), and maximizing it with respect to the parameters, I obtain the maximum likelihood estimates for the population composition and q . This procedure will produce consistent estimates, since the model is identified (there are no observationally equivalent¹⁵ parameters), the parameter

¹⁵Provided that we accept the indeterminacy between behaviorally equivalent outcomes.

space (s_0, s_1, q) is compact and λ is continuous in (s_0, s_1, q) for almost all t ¹⁶.

5.1 Estimation results for the *EF-CH* model

Table 3 presents the parameter values that maximize the log-likelihood function (6) for each possible outcome of the *EF-CH* model. The estimates show the maximum likelihood estimates of the shares of type 0, s_0 , and type 1, s_1 . In the same Table it is also possible to read the estimated continuation probability q .

Table 3: Maximum Likelihood Estimation			
	i / Bi	ii / Bii	$iii / Biii$
-LogL	555.82	620.15	490.06
s_0	0.747** (0.004)	0.361** (0.015)	0.852** (0.014)
s_1	—	0.360** (0.015)	—
q	0.936** (0.029)	0.597** (0.022)	0.283** (0.004)
Obs.	281	281	281
Std. errors in par. ; * sign. 5%; ** sign. at 1%			

The estimated fraction of type 0s in the population is very high in outcomes i / Bi and $iii / Biii$, while it is around thirty per cent in outcome ii / Bii . Note also that the estimates of q vary substantially across outcomes, but in outcome ii / Bii they are close to the uniform randomization value. This last result is suggestive of the presence of a truly randomizing type 0 in the population.

In general, the results in Table 3 suggest that all the outcomes can be estimated using the data. However, these estimates are not informative as regards which of the outcomes is preferable in the sense of being the correct model of behavior and in terms of parsimony of parameters. Regarding this last point, note that outcome ii / Bii needs two parameters to describe the population, s_0 and s_1 , while the other two use only s_0 . This follows from the fact that in these last two outcomes type 1s and type 2s behave in the same way, and hence they are estimated to be the remaining $1 - s_0$ share of the population. To address the issue of which model is preferable, the simplest way is to look at the informational criteria, AIC and BIC, computed for each outcome as shown in Table 4. The information criteria help identifying the correct, most parsimonious model specification. The two criteria are very similar, with the difference that the BIC penalize additional parameters

¹⁶Since the observations are not i.i.d, the convergence in probability of the average log-likelihood to the objective function is guaranteed because the sequence of the average log-likelihood is stochastically equicontinuous.

more severely. In general, for both criteria, the lower the value, the more preferable the model is, in term of parsimony of parameters and explanatory power.

Table 4: Information Criteria for Outcomes' Comparisons

Outcome	Obs.	LogL(null)	LogL(model)	dg. freedom	AIC	BIC
i / Bi	281	.	-555.82	2	1115.64	1122.91
ii / Bii	281	.	-620.15	3	1246.30	1257.22
$iii / Biii$	281	.	-490.06	2	984.13	991.41

The lowest value if the criteria is scored by outcome $iii / Biii$. However, this result may be driven by the fact that the log-likelihood function of outcome $iii / Biii$ is much simpler than those for the other outcomes (see Appendix B). This calls for further investigation, because the information criteria allow for model comparison only and do not provide any diagnostic on the fit of the model. To get a rough measure of how well the outcomes fit the data, I use the maximum likelihood estimates to predict the frequency of termination f_j and the implied take probability p_j for each outcome. First, however, I re-run the analysis considering only first-response data, to gather some more information on the reliability of the estimates.

5.1.1 First-response data

The estimates in Table 3 are obtained by using the full sample of observations. To address possible concerns regarding learning during the game, I here restrict here the analysis to the first response data only, in order to verify if this leads to different results. Table 5 shows the maximum likelihood estimates for the *EF-CH* model

Table 5: Maximum Likelihood Estimation

<i>EF-CH</i> model, first response data			
	i / Bi	ii / Bii	$iii / Biii$
-LogL	50.82	58.51	56.06
s_0	0.890** (0.128)	0.317** (0.052)	0.756** (0.014)
s_1	—	0.291** (0.049)	—
q	0.759** (0.032)	0.586** (0.099)	0.313** (0.005)
Obs.	29	29	29

Std. errors in par. * sign. 5%; ** sign. 1%

The estimates do not change substantially with respect to those in Table 3, suggesting that the data from the full sample can be treated as independent. For the outcomes that

appear in Table 5, is possible to compute the AIC and the BIC, as shown in Table 6

Table 6: Information Criteria for Outcomes' Comparisons

Outcome	Obs.	LogL(null)	LogL(model)	dg. freedom	AIC	BIC
i / Bi	29	.	-50.82	2	105.65	108.39
ii / Bii	29	.	-58.51	3	123.03	127.13
$iii / Biii$	29	.	-56.06	2	116.13	118.87

Now outcome i / Bi is the one that scores the lowest on both criteria. Such a change in the ranking of the three outcomes seems to confirm the suspicion that the low values scored by outcome $iii / Biii$ were an artifact of its extremely simple log-likelihood function.

5.1.2 Fit

I can obtain the distribution of termination as predicted by the outcomes of the *EF-CH* model by substituting the parameters s_0 , s_1 and q into each outcome's \bar{q}_n using the estimates in Table 3. The result of this exercise can be found in Table 7, in which the actual f_j s are reported as well.

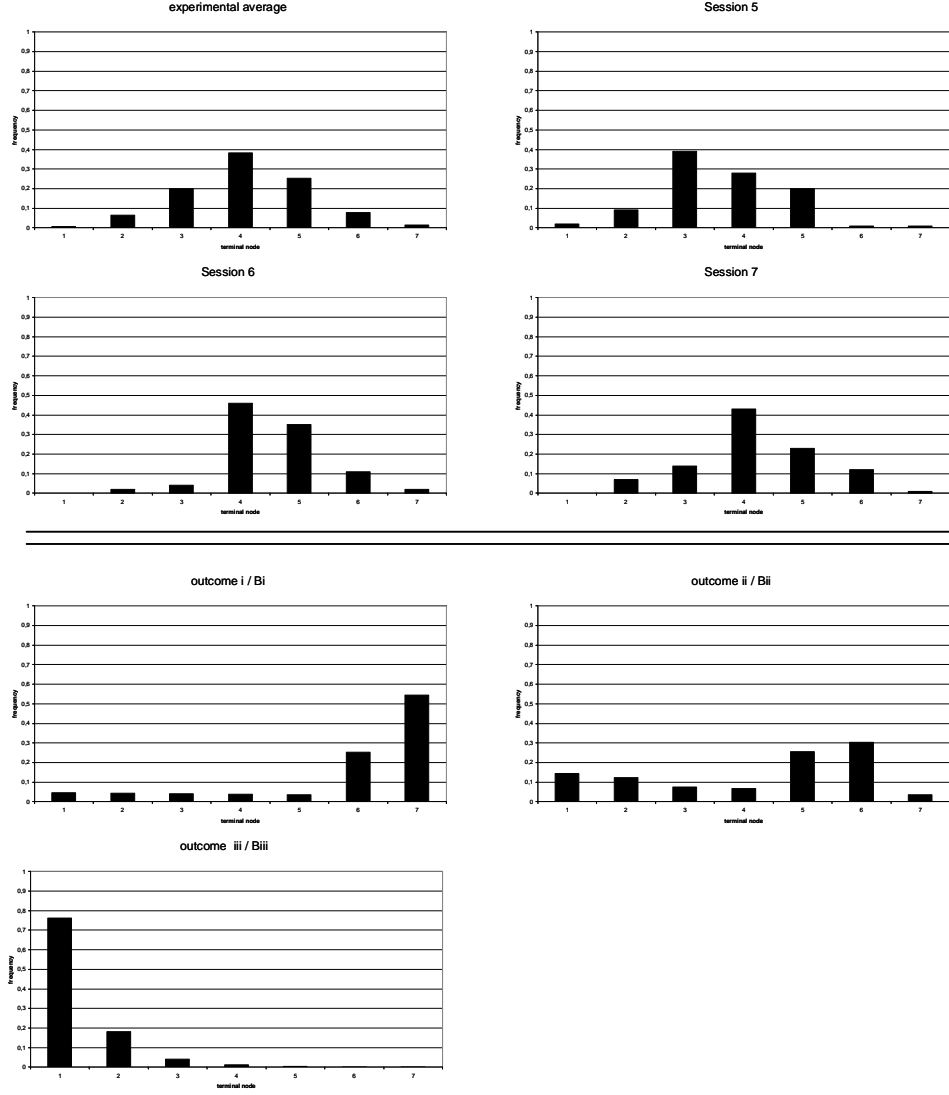
Table 7: Proportion of Observations at each Terminal Node
Empirical and Predicted Distributions

Session	N	f_1	f_2	f_3	f_4	f_5	f_6	f_7
5	100	.02	.09	.39	.28	.20	.01	.01
6	81	.00	.02	.04	.46	.35	.11	.02
7	100	.00	.07	.14	.43	.23	.12	.01
Total 5-7	281	.007	.064	.199	.384	.253	.078	.014
Outcome		\hat{f}_1	\hat{f}_2	\hat{f}_3	\hat{f}_4	\hat{f}_5	\hat{f}_6	\hat{f}_7
i / Bi	—	.045	.043	.040	.039	.036	.253	.544
ii / Bi	—	.144	.123	.074	.066	.255	.303	.035
$iii / Biii$	—	.762	.181	.04	.011	.003	.002	.001

It is easy to see that the predicted terminations occur very early or very late in the game. Outcome $iii / Biii$, which is the one with the lowest informational criteria, predicts an unreasonably high share of terminations the first node. In contrast to that, outcome i / Bi predicts too many terminations at the last two nodes. Outcome ii / Bii seems to fit the empirical termination frequencies reasonably well.

The results in Table 7 can be plotted using an histogram as shown in Figure 6.

Empirical and predicted distributions of termination



The first four panels present the average empirical distribution of termination, together with those of sessions 5, 6 and 7 separately. The laboratory data show a marked tendency to terminate at nodes 3, 4 and 5. The average distribution is almost normal around the games's mean, while the individual sessions show slightly different patterns. Terminations, however, very rarely occur before node 2. The remaining three panels show the predicted terminations' distributions according to the outcomes of the *EF-CH*. Clearly, the outcomes

differ substantially in their predictions. Outcome i / Bi puts too much weight on late terminations with respect to the sample data. Equally far from replicating the empirical pattern, outcome $iii / Biii$ predicts that the largest share of termination takes place at node 1. This leaves us with outcome ii / Bii , which, despite having the highest information criteria among the three, predicts that terminations may take place at all nodes, although largest share of terminations is predicted at nodes 5 and 6.

To investigate further the predictive power of the *EF-CH* outcomes, we can use the data in Table 7 to derive the predicted implied take probability. As discussed in Section 2, empirically such a probability, p_j , increases in the length of the game. This implies that participants are more likely to take the pot at the later nodes, when reached, than earlier in the game. Table 8 presents these figures.

Table 8: Implied take probabilities
Empirical and Predicted Distributions

Session	p_1	p_2	p_3	p_4	p_5	p_6
5	0.02	0.09	0.44	0.56	0.91	0.50
6	0.00	0.02	0.04	0.49	0.72	0.82
7	0.00	0.07	0.15	0.54	0.64	0.92
Total 5-7	0.01	0.06	0.21	0.53	0.73	0.85
	\hat{p}_1	\hat{p}_2	\hat{p}_3	\hat{p}_4	\hat{p}_5	\hat{p}_6
outcome i / Bi	0.05	0.32	0.04	0.05	0.04	0.31
outcome ii / Bii	0.14	0.14	0.10	0.10	0.43	0.89
outcome $iii / Biii$	0.76	0.76	0.70	0.65	0.50	0.66

Again, outcome ii / Bii is the one that follows the empirical data the closest, which is not surprising given that we used the estimated f_j s to derive the empirical take probabilities.

In general, the empirical analysis has shown that the *EF-CH* model outperforms the backward induction results, but falls short in replicating the high number of terminations at the central nodes of the game. Among the possible outcomes, outcome ii / Bii has been found to be the one that better replicates the empirical evidence. In other words, if we assume a population of players composed by the three types in equal shares, and if we assume that the type 0 players randomize with a probability of sixty percent on the choice of continuation, while type 1s and 2s continue to last round (with type 2s terminating there), we can represent fairly well the experimental data. If anything, the cognitive hierarchy approach misses a mechanism to generate terminations at the two central nodes. I will discuss the implication of such a result in section 7.

5.2 Estimation results for the *HEF-CH* case

In what follows, I present the estimation results for the *HEF-CH* outcomes. Note that outcome *I* and *II* are behaviorally indistinguishable from outcome *i* and outcome *ii* respectively; the same holds for outcome *III*⁻ and outcome *iii*. Of the Bayesian extension of the *HEF-CH* case (see Appendix A), only outcome *BIII* prescribes a new behavior. To avoid repetitions, I do not present the estimates for the outcomes that are behaviorally indistinguishable from the ones already discussed. The maximum likelihood estimates are shown in Table 9.

Table 9 : Maximum Likelihood Estimation

	<i>III</i>	<i>IV</i>	<i>BIII</i>	<i>I</i> ⁻	<i>II</i> ⁻
-LogL	576.65	650.29	589.81	635.61	654.22
s_0	0.181** (0.036)	0.077** (0.012)	0.128** (0.007)	0.447 (0.000)	0.479** (0.000)
s_1	0.569** (0.053)	0.975 (0.000)	0.516** (0.003)	0.004** (0.000)	0.000 (0.000)
q	0.792** (0.016)	0.742** (0.028)	0.354** (0.003)	0.370** (0.001)	0.343** (0.000)
Obs.	281	281	281	281	281

Std. errors in parentheses; * sign. 5%; ** sign. at 1%

Recall that in outcomes *III* and *IV* type 1 players continue to the last round, while type 2 players terminate at the central and the first node, respectively. In outcomes *I*⁻ and *II*⁻ type 1 players terminate at the first round, while type 2 players continue to the last round in the former and in the latter terminate at the first node if playing in position *A* and at the last one if in position *B*. This stark difference in the behavior of type 2 players persists in outcome *BIII*, with the difference that type 1 players continue to the last round. With respect to the *EF-CH* case, it is possible to note that the estimates point at more realistic a population composition. On average type 0 players constitute one third of the population. As for q , this varies between thirty and eighty per cent. All the estimates, besides those for outcome *II*⁻ are significant. Table 10 presents the information criteria for these outcomes.

Table 10: Information Criteria for Outcomes' Comparisons

Outcome	Obs.	LogL(null)	LogL(model)	dg. freedom	AIC	BIC
<i>III</i>	281	.	-563.21	3	1159.30	1170.21
<i>IV</i>	281	.	-650.29	3	1304.59	1311.87
<i>BIII</i>	281	.	-589.81	3	1185.63	1196.54
<i>I</i> ⁻	281	.	-635.61	2	1275.22	1282.50
<i>II</i> ⁻	281	.	-638.09	2	1314.45	1325.37

Outcomes *III* and *IV* score the lowest information criteria, noticeably lower than those computed for outcome *i* and outcome *ii*. Before looking at how well these outcomes fit the experimental evidence, I run the same analysis on first-response data only, to check if the results are consistent with the ones just discussed.

5.2.1 First response data

As for the *EF-CH* model, it is possible to use first response data to carry out the maximum likelihood estimation. The results appear in Table 11.

Table 11: Maximum Likelihood Estimation <i>HEF-CH</i> model, first response data			
	<i>III</i>	<i>IV</i>	<i>BIII</i>
-LogL	64.58	69.60	63.51
s_0	0.166** (0.030)	0.935** (0.028)	0.447 (0.000)
s_1	0.723** (0.071)	0.098** (0.024)	0.370** (0.000)
q	0.776** (0.043)	0.769** (0.040)	0.169** (0.000)
Obs.	29	29	29
Std. errors in par. ; * sign. 5%; ** sign. at 1%			

The estimates for outcomes *III*, *IV* and *BIII* do not differ substantially from those in Table 9, obtained by using the full sample. However, the estimation was impossible to carry out for outcome *I*⁻ and *II*⁻, which may be suggestive of a problem either with the outcomes or with the data. Neither of these explanations is troublesome, for two reasons. First, the first response sample consists of 29 observations, which is a small sample to perform maximum likelihood with. Secondly, outcome *I*⁻ and *II*⁻ are those that in the full sample scored the lowest informational criteria, which already then suggested that they were not performing well in explaining the data. For the outcomes in Table 11 we can look at the information criteria.

Table 12: Information Criteria for Outcomes' Comparisons

Outcome	Obs.	LogL(null)	LogL(model)	dg. freedom	AIC	BIC
<i>III</i>	29	.	-64.58	3	135.16	139.26
<i>IV</i>	29	.	-69.60	3	145.21	149.31
<i>BIII</i>	29	.	-63.51	2	131.03	133.77

The AIC and the BIC for outcomes *III* and *BIII* are approximately the same as in the full sample, and outcome *IV* continues to be the least preferable among the alternatives.

This confirms the results obtained above that point at outcome *III* as the preferable among the alternative ones.

5.2.2 Fit

Using the estimates in Table 9, I can obtain the outcomes' predicted distribution of termination. The results of this exercise are presented in Table 13.

Table 13: Proportion of Observations at each Terminal Node
Predicted Distributions

Outcome	\hat{f}_1	\hat{f}_2	\hat{f}_3	\hat{f}_4	\hat{f}_5	\hat{f}_6	\hat{f}_7
<i>III</i>	.120	.105	.303	.185	.041	.117	.129
<i>BIII</i>	.760	.150	.010	.070	.000	.010	.000
<i>IV</i>	.234	.180	.132	.102	.077	.106	.169
<i>I</i> ⁻	.280	.200	.080	.060	.020	.330	.030
<i>II</i> ⁻	.830	.050	.070	.000	.020	.010	.020

Looking at Table 13 is easy to see that the problem of the *EF-CH* model largely persists. In fact, the outcomes fail to predict the large number of terminations at the central round. However, the *HEF-CH* outcomes fare somewhat better than those of the *EF-CH* model. To better understand how, I plot the predicted frequency of termination in a histogram, as can be seen in Figure 7.

The distributions in Figure 7 demonstrate more diverse patterns of terminations than those in Figure 6. In particular, outcome *III*, outcome *IV* and outcome *I*⁻ are doing reasonably well. As before, we can use the predicted frequency of termination to back out the predicted implied take probabilities, which are shown in Table 14.


Table 14: Implied take probabilities
Predicted Distributions

	\hat{p}_1	\hat{p}_2	\hat{p}_3	\hat{p}_4	\hat{p}_5	\hat{p}_6
outcome <i>III</i>	0.12	0.12	0.39	0.39	0.14	0.48
outcome <i>BIII</i>	0.76	0.63	0.11	0.88	0.00	1.00
outcome <i>IV</i>	0.83	0.33	0.58	0.00	0.40	0.33
outcome <i>I</i> ⁻	0.28	0.28	0.15	0.14	0.05	0.92
outcome <i>II</i> ⁻	0.83	0.29	0.58	0.00	0.40	0.33

The number of Table 14 are closer to the actual values than those in Table 6, for the *EF-CH* model. In particular, outcome *III* is the closest to the average empirical data,

with a major difference to it only at node 5.

Predicted distributions of termination



Summing up, the empirical analysis of the *HEF-CH* models suggest that, among the alternatives, outcome III is the outcome that best explains the subjects' behavior when playing the game. Outcome *III* describes a situation in which type 1 players continue to last round, assigning a high probability to the event of the opponent passing the pot back ($\lambda \geq \frac{1}{7}$). Type 2 players, instead, assign to such an event a probability μ strictly lower than that of type 1, and believe that the population consists mostly of type 0 players. Note that, according to the estimates in Table 9, type 2 players have a correct belief on the population composition, but they are mistaken in assigning such a low continuation probability. Since the actual q is estimated at 0.79, instead, it is type 1 players that have the correct belief on q .

6 Discussion

The *EF-CH* model described in section 3 has brought about a variety of findings that merit discussion. Most importantly, the extensive-form cognitive hierarchy assumptions prescribe optimal behaviors that differ from the backward induction equilibrium. In particular, the *HEF-CH* case is such that it is possible for type 2 players to terminate at any node of the game. This result can be interpreted as an interesting way out of the backward induction paradox, as defined by the discrepancy between equilibrium predictions and laboratory behavior. By simply relaxing the equilibrium requirements on players' beliefs it is possible to generate payoff maximizing behaviors that are, in theory, consistent with the experimental data.

However, when brought to the data, the model partially fails the empirical test. While it succeeds in consistently replicating the pattern of the implied take probabilities, it falls short of predicting the almost normal distribution of the frequency of terminations, by assigning a too high probability to the event that the play reaches the final round. There are at least two natural ways to remedy this shortcoming. One possibility is to increase the number of types in the population: higher types, in fact, tend typically¹⁷ to terminate earlier. Alternatively, the mismatch between the predicted and empirical distribution of terminations could be addressed by relaxing the assumption of *constant* randomization from part of type 0 players. The experimental evidence seems to suggest that the central round features some aspects that make it different from the initial and the last one. Probably this follows from some behavioral consideration. Regret, risk aversion, loss aversion most likely come into play and prompt the decision to terminate more at the middle nodes than at the earlier ones.

Both solutions have already been explored, albeit outside the context of the cognitive hierarchy approach, and both have successfully contributed to shed light on players' behavior in the centipede game. McKelvey and Palfrey (1992) introduced altruism as a possible explanatory factor. In particular, they offered a model of incomplete information in which there is the possibility to be matched with an altruistic opponent that always chooses to pass. Allowing for both errors in actions and in beliefs, they show that in equilibrium players adopt mixed strategies that prescribe them to terminate with a probability that increases in the length of the play. According to their estimates, about five per cent of the players are altruists. Compared to this model, the *EF-CH* has the advantage that it does not require the players to be able to derive an equilibrium in mixed strategy. Moreover, it can accommodate the belief of playing with an altruist. However, this comes at the cost of much less sharp predictions.

Kawagoe and Takizawa (2010) instead allow for types higher than type 2 in the context

¹⁷Higher types, in fact, have by definition increasingly sophisticated beliefs about the opponent's behavior. As these beliefs approach the common knowledge assumption, the resulting payoff-maximizing behavior tends to the Nash equilibrium, which, in the context of the Centipede game, is to terminate at the first round.

of the Level k model. When applied to the same game as in Figure 1, they estimate a population composed mostly of type 1s, roughly seventy per cent, and type 2 (the authors assume that type 0s only exist in the minds of type 1s). This is fairly close to the estimates obtained when looking at outcome *III*, where the type 1 players are 57 per cent of the population, type 0s 18 per cent and type 2s are 25 per cent. Such a similarity is encouraging, suggesting that it is reasonable to assume that players' behavior can be explained, in the case of a six-step centipede game, by assuming three types of players only. However, the authors show that for other centipede games the fraction of type 2s and 3s may be sizeable. While it is certainly feasible to try out different population compositions to match the data, I argue that such a fine-tuning fails the spirit of the original research question, namely how far would a realistic model of players' behavior lead us into explaining the centipede game. Admitting type 3 players means to assume that in the experimental population some individuals are prepared to undertake very high cognitive costs to decide how to play. The expensiveness of type 3 reasoning, together with the empirical finding that players' behavior in most games can be described assuming types 0, 1 and 2, speak against the idea of allowing higher types of players.

When compared with other models that address the backward induction paradox, my approach appears substantially different from those put forward in the behavioral literature, which typically invokes factors like altruism, fairness or similar concepts. While the behavioral approach most certainly has some explanatory power over the experimental results that document late termination in the game, it fails to explain the mechanism driving the paradox. This is embed in the payoff structure, that simultaneously makes it attractive to be at the later nodes in the game, where payoffs are higher, and unattractive to reach them, since for every two subsequent nodes the first mover prefers to terminate rather than continue. Altruism, fairness or reciprocity are channels that may explain the preferences to terminate at later stages of the game, where both players are better off, rather than terminate at the beginning of it. Yet, at the last node, reciprocity cannot be enforced, since there are no more moves, and fairness is impossible, since the division of the pot is bound to be uneven by the payoff structure. Altruism is also problematic, since a player, to be altruist, would have to forego at least half of his payoff at the last node. The behavioral solutions appear corollary to the explanation of the paradox, but fail to directly address it. On the contrary, the cognitive hierarchy assumptions offer a mechanism that may reconcile rational play with the continuation patterns found in the data, thus avoiding the paradox and allowing for behavioral refinements.

Finally, the ancillary finding that Bayesian updating does not improve the fit of the model may be relevant to the literature discussing the extent to which Bayesian reasoning really is an appropriate description of human behavior (see, for instance, Cosmides and Tooby, 1996). However, the most likely explanation of this result is the much simplified model of type 1s' beliefs about types 0. As it stands now, the constant randomization assumption clashes with the exponential increase in payoff across nodes, and feels artificial

even granted that it is in the spirit of the model to keep simple type 1s' beliefs about the opponent. A model that allows type 1s to undertake some Bayesian updating might very well fare better than the *EF-CH* specification: to improve the model in this respect is surely another important extension.

7 Conclusions

In this paper, I have extended the cognitive hierarchy model (Camerer et al., 2004) to the class of two-person extensive-form finite games of perfect information. I assumed three types of players: a non-strategic randomizer, type 0, and types 1 and 2, that reason by backward induction according to their beliefs. These were structured according to the conventional cognitive hierarchy hypotheses that type k assumes that the opponent belongs to a mixture of lower-level types, and that type k players, $k \neq 0$, assign the same belief on the behavior of type 0s. Besides studying this model, which I called *EF-CH*, I also studied two extensions. The first allowed type 2s to undertake Bayesian updating on their prior about the opponent's type, and the second, which I called *HEF-CH*, relaxed the assumption that type 1s and 2s have the same belief on the behavior of type 0s.

In both cases, the assumptions have lead to outcomes that improve upon the backward induction result, prescribing, to a different extent, terminations at nodes later than the first. To discriminate which of the possible specifications explains the data best, I estimated and compared them using standard maximum likelihood tools. The empirical analysis produced several results. First, the outcomes of the *EF-CH* model predict terminations either too early or too late in the game to match the data. The best fit is obtained assuming a population composition where the three types are present in even shares, and such that type 0s are more likely to continue than to terminate, while type 1s and 2s continue to the last round. This result can be improved by one of the *HEF-CH* outcomes, assuming a population composed mostly by type 1 players and only in minor part by type 0 players that continue with a high probability. Despite that the fit of this outcome is closer to the data, it still falls short of predicting the large share of terminations that occur during the experiment at nodes 3 and 4. Interestingly, allowing for Bayesian updating does not seem to have any noticeable effect on the model's predictions or empirical performance.

Concluding, the extensive-form cognitive hierarchy model has proven to be reasonably effective and parsimonious in explaining the backward induction paradox. However, the empirical analysis suggests that there is more to the plays of the centipede game than different cognitive levels among the players. Most likely, the experimental behavior is the result of the combined actions of players of different cognitive types and the effect of relevant behavioral factors such as altruism and fairness concerns.

References

- [1] Asheim, G. B. , Dufwenberg, M. , 2003. Deductive reasoning in extensive games. *Econ. J.* 113, 305–325.
- [2] Aumann, R. , 1995. Backward induction and common knowledge of rationality. *Games Econ. Behav.* 8, 6-19.
- [3] Battigalli, P. , 1997. On rationalizability in extensive form games. *J. Econ. Theory.* 74, 40-61.
- [4] Ben-Porath, E. , 1997. Rationality, Nash Equilibrium and Backwards Induction in Perfect Information Games. *Rev. Econ. Stud.* 64, 23-46.
- [5] Brandenburger, A. , Friedenberg A. , 2007. Axioms for backward induction. *Mimeo.*
- [6] Brocas, I. , Camerer, C. , Carillo, J. , Wang, S. , 2009. Measuring attention and strategic behavior in games with private information. CEPR Discussion Paper No. DP7529.
- [7] Brown, A. , Camerer, C. , Lovo, D. , 2009. To review or not to review? Limited strategic thinking at the movie box office. *Mimeo.*
- [8] Camerer, C. , 2003. Behavioral Game Theory: Experiments in strategic interaction. Princeton: Princeton University Press.
- [9] Camerer, C. , Ho, T.H. , Chong, J. K. , 2004. A cognitive hierarchy theory of one-shot games. *Quart. J. Econ.* 119, 861-898.
- [10] Cosmides, L. , Tooby, J. , 1996. Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty. *Cognition.* 58, 1-73.
- [11] Costa-Gomes, M. , Crawford, V. , Broseta, B. , 2001. Cognition and behavior in normal-form games: An experimental study. *Econometrica.* 69, 1193-1235.
- [12] Costa-Gomes, M. , Crawford, V. , 2006. Cognition and behavior in two-person guessing games: An experimental study. *Amer. Econ. Rev.* 96, 1737-1768.
- [13] Costa-Gomes, M. , Crawford, V. , Iriberri, N. , 2009. Comparing models of strategic thinking in van Huyck, Battalio, and Beil's Coordination Games. *JEEA.* 7, 365–376.
- [14] Crawford, V. , Iriberri, N. , 2007a. Level-k auctions: can a non-equilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica.* 75, 1721–1770.

- [15] Crawford, V. , Iriberri, N. , 2007b. Fatal attraction: Salience, naivete, and sophistication in experimental hide-and-seek games. *Amer. Econ. Rev.* 97, 1731-1750.
- [16] Ellingsen T. , Östling, R. , 2006. When Does Communication Improve Coordination? *Amer. Econ. Rev.*, forthcoming.
- [17] Fey M. , McKelvey, R. , Palfrey, T.R. , 1996. An experimental study of constant sum Centipede games. *Int. J. Game Theory.* 25, 269-287.
- [18] Fudenberg, D. , Levine, D. K. , 1998. *The Theory of Learning in Games*. Cambridge: MIT Press.
- [19] Gerber, A. , Wichardt, P.C. , 2010. Iterated reasoning and welfare-enhancing instruments in the centipede game. *J. Econ. Behav. Organ.* 74, 123-136.
- [20] Hart, S. , 2002. Evolutionary dynamics and backward induction. *Games Econ. Behav.* 41, 227-264.
- [21] Ho, T. H. , Camerer, C. , Weigelt, K. , 1998. Iterated dominance and iterated best response in experimental ‘p-beauty contests’. *Amer. Econ. Rev.* 39, 649-660.
- [22] Ho, T. H. , Wang, X. , Camerer, C. , 2008. Individual differences in EWA learning with partial payoff information. *Econ. J.* 118, 37-59.
- [23] Kawagoe, T. , Takizawa, H. , 2010. Level- k Analysis of Experimental Centipede Games. Available at SSRN: <http://ssrn.com/abstract=1289514>.
- [24] Kreps D. M. , Wilson, R. , 1982. Sequential Equilibria. *Econometrica.* 50, 863-894.
- [25] Kreps D. M. , Milgrom, P. , Roberts, J. , Wilson, R. , 1982. Rational Cooperation in the finitely repeated Prisoner’s Dilemma. *J. Econ. Theory.* 27, 245-52.
- [26] Kuhn, H. W. , 1950. Extensive games. *Proc. Natl. Acad. Sci. U.S.A.* 36, 570-576.
- [27] Kuhn, H. W. , 1953. *Classics in Game Theory*. Princeton: Princeton University Press.
- [28] List, J. , Levitt, S. , Sadoff, S. , 2010. Checkmate: Exploring Backward Induction Among Chess Players Checkmate: exploring backward induction among chess players. *Amer. Econ. Rev.*, forthcoming.
- [29] McKelvey, R. , Palfrey, T. , 1992. An experimental study of the centipede game. *Econometrica:* 60, 803-836.
- [30] McKelvey, R. , Palfrey, T. , 1995. Quantal response equilibria for normal-form games. *Games Econ. Behav.* 10, 6-38.

- [31] McKelvey, R. , Palfrey, T. , 1998. Quantal response equilibria for extensive form games. *Exper. Econ.* 1, 9-41.
- [32] Nagel, R. , 1993. Experimental results on interactive competitive guessing. Discussion paper No. B-236, University of Bonn.
- [33] Nagel, R. , 1995. Unraveling in guessing games: An experimental study. *Amer. Econ. Rev.* 85, 1313-1326.
- [34] Nagel, R. , Tang, F. , 1998. Experimental Results on the Centipede Game in Normal Form: An Investigation on Learning. *J. Math. Psychol.* 42, 356-384.
- [35] Nöldeke G. , Samulsson, L. , 1997. A Dynamic Model of Equilibrium Selection in Signaling Markets. *J. Econ. Theory.* 73, 118-156.
- [36] Östling, R. , Wang, T. , Chou, E. , Camerer, C. , 2009. Testing Game Theory in the field: swedish LUPI lottery games. SSE/EFI Working Paper Series in Economics and Finance No. 671, Stockholm School of Economics.
- [37] Palacios-Huerta, I. , Volij, O. , 2009. Field Centipedes. *Amer. Econ. Rev.* 83, 1281-1302.
- [38] Rapoport, A. , Stein, W. E. , Parco, J. E. , Nicholas, T. E. , 2003. Equilibrium Play and Adaptive Learning in a Three-Person Centipede Game. *Games Econ. Behav.* 43, 239-265.
- [39] Reny, P. , 1998. Rationality, Common Knowledge, and the Theory of Games. Technical Report, University of Western Ontario.
- [40] Rosenthal, R. , 1982. Games of perfect information, predatory pricing, and chainstore paradox. *J. Econ. Theory.* 25, 92-100.
- [41] Selten, R. , 1965. Spieltheoretische Behandlung eines Oligopolmodells mit nachfrage-tragheit. *Z. Staatswissenschaft* 121, 301- 324.
- [42] Selten, R. , 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. *Internat. J. Game Theory.* 4, 25-55.
- [43] Stahl, D. , Wilson, P. W. , 1994. Experimental evidence on players' models of other players. *J. Econ. Behav. Organ.* 25, 309-327.
- [44] Stahl, D. , Wilson, P. W. , 1995. On Players' Models of Other Players: Theory and Experimental Evidence. *Games Econ. Behav.* 10, 218-254.

8 Appendix A: Bayesian updating

Equation A.1 shows the payoff-maximizing strategies for type 2 players in the Bayesian *EF-CH* model:

$$\begin{aligned} t_A^* &= \begin{cases} c_1 c_3 c_5 & \text{if } \sigma_0 \geq \tilde{\lambda} \text{ and } \lambda \geq \frac{1}{7} \\ c_1 c_3 t_5 & \text{if } \sigma_0 < \tilde{\lambda} \text{ and } \lambda \geq \frac{1}{7} \\ t_1 t_3 t_5 & \text{otherwise} \end{cases} \\ t_B^* &= \begin{cases} c_2 c_4 t_6 & \text{if } \lambda \geq \frac{1}{7} \\ t_2 t_4 t_6 & \text{otherwise} \end{cases} \end{aligned} \quad (\text{A.1})$$

where $\tilde{\lambda} = \frac{1}{7\lambda^3 - \lambda^2 + 1}$.

We can now look at the effect of allowing for Bayesian updating in the *HEF-CH* model. An informal argument may illustrate how Bayesian updating would affect the outcomes in the game. Consider a player i of type 2 at the beginning of the game. Suppose his beliefs fall in outcome III, meaning that at the beginning of the game it is optimal for i to terminate at node 3 (or 4 if he plays in role B). With Bayesian updating, however, once reached the second round of play, i may decide to continue further. This happens because i knows that only type 0s terminate before the last round, since by assumption he knows λ . Thus, upon receiving the pot back, i revises downward the probability of being playing with a type 0. In particular, player i may revise σ_0 downward enough to decide to continue rather than stop at node 3 (or 4). Looking at Figure 2, this would correspond to a drop in σ_0 from outcome III to outcome II.

As before, I will consider separately the two cases $\lambda \geq \frac{1}{7}$ and $\lambda < \frac{1}{7}$, and start with the former. The behavior of type 1 players' remains the same as that described in equation 1, while the payoff maximizing strategies for type 2s do not. t^* can still be derived by backward induction, now making use at each node of the updated beliefs:

$$\begin{aligned} t_A^* &= \begin{cases} c_1 c_3 c_5 & \text{if } \sigma_0 > \frac{1}{7\mu_0^3 - \mu_0^2 + 1} \\ c_1 c_3 t_5 & \text{if } \sigma_0 < \frac{1}{7\mu_0^3 - \mu_0^2 + 1} \\ t_1 t_3 t_5 & \text{if } \sigma_0 > 0.93 \text{ and } \mu_0 < \frac{1}{7} \end{cases} \\ t_B^* &= c_2 c_4 t_6 \end{aligned} \quad (\text{A.2})$$

A type 2 player in role B has only one payoff maximizing strategy, continuing to the last node: the fact that the opponent passes at node one conveys enough information to make it optimal for the type 2 players to continue the game to the last node. Figure A1

illustrates the (t_A^*, t_B^*) strategies in the $(\mu_0; \sigma_0)$ plane.




Fig. A1. Outcomes of the HEF-CH Bayesian case

Outcome BI is defined by the beliefs that make it payoff maximizing to play $t_A^* = c_1 c_3 c_5$ and $t_B^* = c_2 c_4 t_6$. In outcome BII it is $t_A^* = c_1 c_3 t_5$, $t_B^* = c_2 c_4 t_6$ and in outcome BIII $t_A^* = t_1 t_3 t_5$ and $t_B^* = c_2 c_4 t_6$. The termination probabilities are given in Appendix B.

This case presents two main differences with respect to the non-Bayesian *EF-CH*. First, the behavior $t_A^* = c_1 t_3 t_5$, $t_B^* = c_2 t_4 t_6$ is no more payoff maximizing (in expectation) for any beliefs. In fact, a player starting off with the beliefs that made it payoff maximizing to stop at node 3 (outcome III), once at node 3 revises σ_0 downward enough to keep playing at least one more node. Secondly, continuation at node 5 is now payoff maximizing for a smaller set of beliefs than for the *EF-CH* case. When node 5 is reached, two opposing considerations come into play. On the one hand, a player at node 5 has to believe in a high σ_0 to choose to continue, because c_5 is profitable only with a type 0 opponent. On the other, however, the likelihood of being playing with a type 0 has been revised downwardly, given that the game has not been interrupted. The beliefs for which it is optimal to play c_5 , then, shrink with respect to the *EF-CH* model.

In case $\lambda < \frac{1}{7}$, the Bayesian updating takes place only in the first round and only for those type 2s that are playing in position B , since those in the other role terminate at node 1, as shown in equation ???. Knowing that all type 1s go out after the first round, those type 2s that receive the pot at node 2 revise σ_0 to $\sigma_0^{R1} = 1$. This revision affects type 2s' behavior differently for different μ_0 . If $\mu_0 \geq \frac{1}{7}$, then type 2s terminate at the last round, otherwise they take the pot immediately.

9 Appendix B: termination probabilities \bar{q}_n

Remark 1 *Procedure to find the payoff maximizing strategies in the case $\rho_1 = \rho_2 = \frac{1}{2}$.*

Consider a type 1 player in role A. If $\rho_1 = \frac{1}{2}$, at the last node $E\pi(c_5) = \frac{1}{2}\pi(c_6) + \frac{1}{2}\pi(t_6) = 14.4$ and $\pi(t_5) = 6.4$, meaning that his optimal choice is c_5 . Reasoning backwardly, one step ahead, the expected value of playing c_3 is $E\pi(c_3) = \frac{1}{2}\pi(c_4) + \frac{1}{2}\pi(t_4) = 7.6$, greater than $\pi(t_3) = 1.6$. At the initial node, again it is the case that $E\pi(c_1) = \frac{1}{2}\pi(c_2) + \frac{1}{2}\pi(t_2) = 4 > 0.4 = \pi(t_1)$. Then, type 1's optimal strategy is $r_A^* = c_1c_3c_5$. When in role B, a type 1 plays t_6 at the last node. At the fourth, $E\pi(c_4) = \frac{1}{2}\pi(c_5) + \frac{1}{2}\pi(t_5) = 7.2 > 3.2 = E\pi(t_4)$. At node two, the value of terminating is $\pi(t_2) = 0.8$ and that of continuing is $E\pi(c_2) = \frac{1}{2}\pi(c_3) + \frac{1}{2}\pi(t_3) = 3.8$. Summing up, the payoff maximizing strategies for a type 1 type are: $r_A^* = c_1c_3c_5$ and $r_B^* = c_2c_4t_6$. Reasoning in the same way for a type 2 player, we have that $t_A^* = c_1c_3t_5$ and $t_B^* = c_2c_4t_6$, which is the best reply to (r_A^*, r_B^*) .

Termination probabilities outcomes i / I / BI

$$\begin{aligned}\bar{q}_1 &= s_0(1 - q) \\ \bar{q}_2 &= s_0^2 q(1 - q) + (s_0 - s_0^2)(1 - q) \\ \bar{q}_3 &= s_0^2 q^2(1 - q) + (s_0 - s_0^2)q(1 - q) \\ \bar{q}_4 &= s_0^2 q^3(1 - q) + s_0 q - s_0^2 q - s_0 q^2 + s_0^2 q^2 \\ \bar{q}_5 &= s_0^2 q^4(1 - q) + s_0 q^2 - s_0 q^3 - s_0^2 q^2 + s_0^2 q^3 \\ \bar{q}_6 &= s_0^2 q^5(1 - q) + s_0 q^2 - s_0^2 q^2 + 1 - 2s_0 + s_0^2 \\ \bar{q}_7 &= s_0^2 q^6 + s_0 q^3 - s_0^2 q^3\end{aligned}$$

Termination probabilities outcomes ii / II / BII

$$\begin{aligned}\bar{q}_1 &= s_0(1 - q) \\ \bar{q}_2 &= s_0(1 - q) - (1 - q)^2 s_0^2 \\ \bar{q}_3 &= s_0^2 q^2 - s_0^2 q^3 + s_0 q - s_0 q^2 - s_0^2 q + s_0^2 q^2 \\ \bar{q}_4 &= s_0^2 q^3 - s_0^2 q^4 + s_0 q - s_0 q^2 - s_0^2 q + s_0^2 q^2 \\ \bar{q}_5 &= s_0^2 q^4 - s_0^2 q^5 + s_0 q^2 - s_0 q^3 - s_0^2 q^2 + s_0^2 q^3 + (1 - s_0 - s_1)(s_0 q^2 + 1 - s_0) \\ \bar{q}_6 &= s_0^2 q^5 - s_0^2 q^6 + s_0 q^3 - s_0^2 q^3 + s_0 s_1 q^2(1 - q) + s_1 - s_0 s_1 \\ \bar{q}_7 &= s_0 q^3(s_0 q^3 + s_1)\end{aligned}$$

Termination probabilities outcomes iii / III*

$$\begin{aligned}\bar{q}_1 &= 1 - s_0 q \\ \bar{q}_2 &= s_0 q - s_0^2 q^2 \\ \bar{q}_3 &= s_0^2 q^2(1 - q) \\ \bar{q}_4 &= s_0^2 q^3(1 - q) \\ \bar{q}_5 &= s_0^2 q^4(1 - q) \\ \bar{q}_6 &= s_0^2 q^5(1 - q) \\ \bar{q}_7 &= s_0^2 q^6\end{aligned}$$

Termination probabilities outcome III

$$\begin{aligned}
\bar{q}_1 &= s_0(1 - q) \\
\bar{q}_2 &= s_0^2 q(1 - q) + (s_0 - s_0^2)(1 - q) \\
\bar{q}_3 &= s_0^2 q^2(1 - q) + (s_0 - s_0^2)q(1 - q) + (1 - s_0 - s_1)s_0 q + (1 - s_0 - s_1)(1 - s_0) \\
\bar{q}_4 &= s_0^2 q^3(1 - q) + s_0 q^2 - s_0^2 q^2 - s_0 s_1 q^2 + s_0 s_1 q - s_0 s_1 q^2 + s_1 - s_1^2 - s_1 s_0 \\
\bar{q}_5 &= s_0^2 q^4(1 - q) + s_1 s_0 q^2 - s_1 s_0 q^3 \\
\bar{q}_6 &= s_0^2 q^5(1 - q) + s_1 s_0 q^2 + s_1^2 \\
\bar{q}_7 &= s_0^2 q^6 + s_1 s_0 q^3
\end{aligned}$$

Termination probabilities outcome IV

$$\begin{aligned}
\bar{q}_1 &= s_0(1 - q) + 1 - s_1 - s_0 \\
\bar{q}_2 &= s_0^2 q(1 - q) + s_0(1 - s_1 - s_0)q + s_0 s_1(1 - q) + (1 - s_0 - s_1)s_1 \\
\bar{q}_3 &= s_0^2 q^2(1 - q) + s_0 s_1 q(1 - q) \\
\bar{q}_4 &= s_0^2 q^3(1 - q) + s_0 s_1 q^2(1 - q) \\
\bar{q}_5 &= s_0^2 q^4(1 - q) + s_0 s_1 q^3(1 - q) \\
\bar{q}_6 &= s_0^2 q^5(1 - q) + s_0 s_1 q^4 + s_1^2 \\
\bar{q}_7 &= s_0^2 q^6 + s_1 s_0 q^3
\end{aligned}$$

Termination probabilities outcome BIII

$$\begin{aligned}
\bar{q}_1 &= s_0(1 - q) + s_1 \\
\bar{q}_2 &= s_0^2 q(1 - q) + s_0 s_1 q + s_0(1 - s_0 - s_1)(1 - q) + s_1(1 - s_0 - s_1) \\
\bar{q}_3 &= s_0^2 q^2(1 - q) + s_0(1 - s_0 - s_1)q(1 - q) \\
\bar{q}_4 &= s_0^2 q^3(1 - q) + s_0(1 - s_0 - s_1)q^2(1 - q) \\
\bar{q}_5 &= s_0^2 q^4(1 - q) + s_0(1 - s_0 - s_1)q^3(1 - q) \\
\bar{q}_6 &= s_0^2 q^5(1 - q) + s_0(1 - s_0 - s_1)q^4(1 - q) + s_0(1 - s_0 - s_1)q^3 + (1 - s_0 - s_1)^2 \\
\bar{q}_7 &= s_0^2 q^6 + s_0(1 - s_0 - s_1)q^3
\end{aligned}$$

Termination probabilities outcome I*

$$\begin{aligned}
\bar{q}_1 &= s_1 + s_0(1 - q) \\
\bar{q}_2 &= s_0^2 q(1 - q) + s_0 s_1 q + s_0(1 - s_0 - s_1)(1 - q) + s_1(1 - s_0 - s_1) \\
\bar{q}_3 &= s_0^2 q^2(1 - q) + s_0(1 - s_0 - s_1)q(1 - q) \\
\bar{q}_4 &= s_0^2 q^3(1 - q) + s_0(1 - s_0 - s_1)q^2(1 - q) \\
\bar{q}_5 &= s_0^2 q^4(1 - q) + s_0(1 - s_0 - s_1)q^3(1 - q) \\
\bar{q}_6 &= s_0^2 q^5(1 - q) + s_0(1 - s_0 - s_1)q^4(1 - q) + s_0(1 - s_0 - s_1)q^3 + (1 - s_0 - s_1)^2 \\
\bar{q}_7 &= s_0^2 q^6 + s_0(1 - s_0 - s_1)q^3
\end{aligned}$$

Termination probabilities outcome II*

$$\begin{aligned}
\bar{q}_1 &= 1 - s_0 q \\
\bar{q}_2 &= s_0^2 q(1 - q) + s_0 s_1 q
\end{aligned}$$

$$\begin{aligned}
\bar{q}_3 &= s_0^2 q^2 (1 - q) + s_0 (1 - s_0 - s_1) q (1 - q) \\
\bar{q}_4 &= s_0^2 q^3 (1 - q) \\
\bar{q}_5 &= s_0^2 q^4 (1 - q) + s_0 (1 - s_0 - s_1) q^2 (1 - q) \\
\bar{q}_6 &= s_0^2 q^5 (1 - q) + s_0 (1 - s_0 - s_1) q^3 \\
\bar{q}_7 &= s_0^2 q^6
\end{aligned}$$

Termination probabilities outcome 4s-I

$$\begin{aligned}
\bar{q}_1 &= s_0^2 (1 - q) + (s_0 - s_0^2) (1 - q) \\
\bar{q}_2 &= s_0^2 q (1 - q) + (s_0 - s_0^2) (1 - q) \\
\bar{q}_3 &= s_0^2 q^2 (1 - q) + (s_0 - s_0^2) q (1 - q) \\
\bar{q}_4 &= s_0^2 q^3 (1 - q) + (s_0 - s_0^2) q^2 + (s_0 - s_0^2) q (1 - q) + s_1^2 \\
&\quad + s_1 (1 - s_0 - s_1) + (1 - s_0 - s_1)^2 + (1 - s_0 - s_1) s_1 \\
\bar{q}_5 &= s_0^2 q^4 + (s_0 - s_0^2) q^2
\end{aligned}$$

Termination probabilities outcome 4s-II

$$\begin{aligned}
\bar{q}_1 &= s_0^2 (1 - q) + (s_0 - s_0^2) (1 - q) \\
\bar{q}_2 &= s_0^2 q (1 - q) + s_0 s_1 (1 - q) + s_0 (1 - s_0 - s_1) (1 - q) \\
\bar{q}_3 &= s_0^2 q^2 (1 - q) + (s_0 - s_0^2) q (1 - q) + s_0 q (1 - s_0 - s_1) + (1 - s_0 - s_1) (1 - s_0) \\
\bar{q}_4 &= s_0^2 q^3 (1 - q) + (s_0 - s_0^2) q^2 + s_0 s_1 q (1 - q) + s_1^2 + s_1 (1 - s_0 - s_1) \\
\bar{q}_5 &= s_0^2 q^4 + s_0 s_1 q^2
\end{aligned}$$

Termination probabilities outcome 4s-III

$$\begin{aligned}
\bar{q}_1 &= s_0^2 (1 - q) + (s_0 - s_0^2) (1 - q) + (1 - s_1 - s_0) \\
\bar{q}_2 &= s_0^2 q (1 - q) + s_0 s_1 (1 - q) \\
\bar{q}_3 &= s_0^2 q^2 (1 - q) + (s_0 - s_0^2) q (1 - q) \\
\bar{q}_4 &= s_0^2 q^3 (1 - q) + (s_0 - s_0^2) q^2 + s_0 s_1 q (1 - q) + s_1^2 + s_1 (1 - s_0 - s_1) \\
\bar{q}_5 &= s_0^2 q^4 + s_0 s_1 q^2
\end{aligned}$$

Termination probabilities outcome 4s-IV

$$\begin{aligned}
\bar{q}_1 &= s_0^2 (1 - q) + (s_0 - s_0^2) (1 - q) + (1 - s_1 - s_0) \\
\bar{q}_2 &= s_0^2 q (1 - q) + s_0 q (1 - s_0 - s_1) + s_0 s_1 (1 - q) + s_1 (1 - s_0 - s_1) \\
\bar{q}_3 &= s_0^2 q^2 (1 - q) + s_0 s_1 q (1 - q) \\
\bar{q}_4 &= s_0^2 q^3 (1 - q) + s_0 s_1 q^2 + s_1^2 + s_0 s_1 q (1 - q) \\
\bar{q}_5 &= s_0^2 q^4 + s_0 s_1 q^2
\end{aligned}$$

Termination probabilities outcome 4s-I⁻

$$\begin{aligned}
\bar{q}_1 &= s_0^2 (1 - q) + (s_0 - s_0^2) (1 - q) + s_1 \\
\bar{q}_2 &= s_0^2 q (1 - q) + s_0 s_1 q + s_0 (1 - s_0 - s_1) (1 - q) + s_1 (1 - s_0 - s_1) \\
\bar{q}_3 &= s_0^2 q^2 (1 - q) + s_0 (1 - s_0 - s_1) q (1 - q) \\
\bar{q}_4 &= s_0^2 q^3 (1 - q) + s_0 (1 - s_0 - s_1) q^2 + s_0 (1 - s_0 - s_1) q (1 - q) + (1 - s_0 - s_1)^2 \\
\bar{q}_5 &= s_0^2 q^4 + s_0 q^2 (1 - s_0 - s_1)
\end{aligned}$$

Termination probabilities outcome 4s-II⁻

$$\bar{q}_1 = s_0^2 (1 - q) + (s_0 - s_0^2)(1 - q) + 1 - s_0$$

$$\bar{q}_2 = s_0^2 q(1 - q) + s_0 s_1 q$$

$$\bar{q}_3 = s_0^2 q^2(1 - q) + s_0(1 - s_0 - s_1)q(1 - q)$$

$$\bar{q}_4 = s_0^2 q^3(1 - q) + s_0(1 - s_0 - s_1)q^2$$

$$\bar{q}_5 = s_0^2 q^4$$

Termination probabilities outcome 4s-III⁻

$$\bar{q}_1 = s_0^2 (1 - q) + (s_0 - s_0^2)(1 - q) + 1 - s_0$$

$$\bar{q}_2 = s_0^2 q(1 - q) + (s_0 - s_0^2)q$$

$$\bar{q}_3 = s_0^2 q^2(1 - q)$$

$$\bar{q}_4 = s_0^2 q^3(1 - q)$$

$$\bar{q}_5 = s_0^2 q^4$$