Working Paper No. 322

Finite-Order Beliefs and Welfare-Enhancing Instruments in the Centipede Game

Anke Gerber and Philipp C. Wichardt

December 2007

Finite-Order Beliefs and Welfare-Enhancing Instruments in the Centipede Game

Anke Gerber[†]

Department of Economics, University of Hamburg

Philipp C. Wichardt[‡]

Department of Economics, University of Bonn

This Version: December 19, 2007

Abstract This paper investigates the effectiveness of two instruments designed to defer termination in the centipede game: an insurance against termination by the opponent, and an option to offer the opponent a bonus for not terminating the game. The rational prediction in both cases is passing until close to the end. Empirically, however, only the bonus option is used by the subjects. The results indicate that subjects readily understand the strategic effect of the bonus, which, once offered, renders passing until close to the end the strictly dominant strategy for both players. Yet, they fail to realise the slightly more involved strategic signal entailed in the insurance, namely that passing until close to the end is a strictly dominant strategy for an insured player. In order to further investigate this effect, we propose a simple behavioural model based on level-k thinking and show that it is largely consistent with the data.

Keywords: Centipede Game, Cooperation, Level-k Thinking
JEL code: C91, D61

[†]Postal address for correspondence: Department of Economics, University of Hamburg, Von-Melle-Park 5, D-20146 Hamburg, Germany; e-mail: anke.gerber@wiso.uni-hamburg.de.
[‡]e-mail: philipp.wichardt@uni-bonn.de.

1 Introduction

Many economic as well as social interactions bear characteristics of a repeated trust game in that both aggregate and individual welfare can be considerably increased through repeated (mutual) transfers. Yet, as transfers often are not enforceable, any such contribution entails the risk of not being reciprocated. Joint ventures of firms, for example, are typically characterized by repeated investments or transfers of technological know-how in order to increase the productivity of the other party. However, at any point in time this leads to an unequal split of the benefits and there is a permanent risk that transfers are not reciprocated once they have been obtained. Co-authorship is another example, in particular if authors alternate in working on a paper. Any effort made by one author further improves the paper and, hence, benefits both authors. Yet, only one currently bears the cost.

A convenient strategic representation of such interactions is the centipede game (Rosenthal, 1981). In this game, players alternately choose whether to terminate the interaction or to pass over to the opponent. Passing always increases aggregate payoffs. However, if the game is terminated, joint payoffs are split unequally with the terminating player getting the larger share of the pie. Since the game is finite, a standard backwards induction argument yields immediate termination as the unique Nash equilibrium outcome. However, neither immediate termination nor passing until close to the end is typically observed in experimental studies of the centipede game. Rather, in the majority of cases the game is terminated after about 2/3 of the maximum number of rounds (e.g. McKelvey and Palfrey, 1992; Nagel and Tang, 1998). Thus, from the perspective of a social planner, considerable welfare gains are lost, even though the loss is smaller than predicted by standard theory.

Inspired by these observations, we set out to investigate empirically the effectiveness of two instruments designed to increase cooperation, i.e. defer termination, in the centipede game. The first instrument is a costly insurance against the loss due to termination by the opponent. The purchase of the insurance is observed by both players prior to the play of the game. The second instrument is a bonus which players can offer to their opponent (in terms of a binding promise) for not terminating the game. Both instruments are designed so as to — the-

oretically — induce passing at least up to the last or next but last node, and independent of the actual length of the game.

Obviously, the two instruments differ significantly in the way they induce players to abstain from early termination. The insurance, on the one hand, banks on a signalling effect: buying the insurance is like credibly announcing that one is not afraid of termination by the opponent and, therefore, will always pass.[1] If the opponent understands this signal, passing up to his/her last decision node becomes optimal. The bonus, on the other hand, simply rewards the opponent for permanent passing and renders it the strictly dominant action for both the offering player and the recipient of the offer. In both cases, passing until (close to) the end is the unique Nash equilibrium outcome if the respective option is taken by at least one of the players. Therefore, in any pure strategy subgame perfect equilibrium of the extended game, exactly one player should buy the insurance or offer the bonus, respectively.[2]

In the laboratory experiment, however, the insurance option fails badly. Only few subjects insure themselves against termination by the opponent and aggregate observed play is not significantly different from the standard centipede game without insurance. By contrast, the opportunities offered by the bonus option are readily understood and seized. For the vast majority of pairs at least one of the players offers the bonus when this option is available. In fact, in most cases the bonus is offered by the first mover, in accordance with the unique subgame perfect Nash equilibrium for this game. Accordingly, observed termination is significantly delayed as compared to the standard centipede game without bonus.

On a general level, the results demonstrate that mechanisms which require players to recognize that their opponents will not play a strictly dominated strategy (as in case of the bonus) are more effective than mechanisms that rely on an iteration of this reasoning (as in case of the insurance). This is consistent with earlier findings showing that subjects avoid dominated strategies but do not necessarily iterate the dominance reasoning (e.g. Katok, Sefton, and Yavas,

---

[1]The only exception to this is the last node, where the insurance has no effect.

[2]In case of the bonus option, there is a unique subgame perfect Nash equilibrium, in which only the first mover offers the bonus. In case of the insurance, there also is a mixed equilibrium in which both players purchase the insurance with positive probability.

2002).[3]

However, it is also interesting to see whether any further light can be shed on the specific reasons for the failure of the insurance as compared to the bonus, especially in view of iterated dominance arguments. In order to do so, we propose a simple behavioural model which is based on the assumption that subjects imperfectly analyse the strategic structure of the centipede game and that the degree of rationality varies among subjects. The type of model considered is commonly referred to as level-k thinking and is usually used to study initial responses in normal form games (cf. Stahl and Wilson, 1994 and 1995; Nagel, 1995; Ho, Camerer, and Weigelt, 1998; Costa-Gomes, Crawford, and Broseta, 2001; Crawford and Iriberri, 2007a and 2007b).

Level-k-thinking models are based on the assumption that more sophisticated, level-k players best respond to the behaviour of level-$(k-1)$ players, where the iteration process is anchored in an assumed non-strategic behaviour of the level-0 type, L0. Concerning the behaviour of L0, most approaches assume a uniform distribution over all available strategies (e.g. Stahl and Wilson, 1995; Crawford and Iriberri, 2007a). Such a specification is appealing for normal form games as the "value" of any strategy in these games is hard to assess without thinking strategically (with the potential exception of strictly dominated strategies). But also more refined specifications are used, including (partly) estimated distributions over available strategies (Ho, Camerer, and Weigelt, 1998) as well as adjustments in view of the salience of certain actions (Crawford and Iriberri, 2007b). So, as Crawford and Iriberri (2007a, p. 1734) point out, eventually *"the best specification of L0 is an empirical question"*.

For the purpose of the current analysis, we assume that the L0-types always play pass. In case of the centipede game this assumption seems reasonable since all pass-take decisions appear to be similar if a player does not employ any strategic considerations about her opponent's behaviour. Hence, as take is costly in

---

[3] Arguably, also framing effects (c.f. Andreoni, 1995) may have had their share in performance difference between the two mechanisms. In particular, the bonus may have been associated with a positive and the insurance with negative or defensive connotations with the respective positive (bonus) and negative (insurance) prospects of reciprocation and, accordingly, consequences for usage frequencies. However, as subjects' responses to the ex post questionnaire are not indicative of such effects, we consider them to be comparably weak.

terms of forgone payoffs, a non-strategic player, L0, can be expected to play pass at all her decision nodes.

The above "all-pass" assumption for L0 is also convenient in view of the specific extensive-form structure of the centipede game as it ties the level of thinking to the steps of iterated backward reasoning. Thus, we can easily differentiate different levels of iterated backward reasoning within the model. As we aim to scrutinise the driving forces behind the performance differences between the two treatment specifications — the insurance and the bonus — this is preferable, for example, to a crazy types approach (Kreps et al., 1982) as used by McKelvey and Palfrey (1992), which presumes the existence of fully rational and crazy types only. In particular, with rational and crazy types only, we would expect either everyone or no one to purchase the insurance, depending on the behaviour associated with crazy types. Yet, despite an admittedly infrequent usage of the insurance, this is not what we observe.

By contrast, observed behaviour in all treatments is largely consistent with our simple belief based model if, based on data obtained in the standard treatment, we assume that the majority of agents best-responds to a belief of order at most 1. In that case, only the benefit of the bonus but not the signalling effect of the insurance can be understood by the majority of the subjects. In fact, in order to understand the effect of the insurance, a player has to notice that buying the insurance changes the opponent's best-response behaviour from take to pass at least once. This, however, necessitates that the player also considers the opponent's awareness of a change in the player's own best-response behaviour due to an insurance purchase.[4] Thus, it requires at least 2nd-order beliefs, $k = 2$, which only a minority of subjects seem to hold. By contrast, in the bonus treatment it is evident for all players with at least 1st-order beliefs that passing is a dominant strategy for a player who was offered a bonus. Accordingly, given that few subjects seem to have more than 1st-order beliefs in the standard treatment, we expect infrequent usage of the insurance and frequent usage of the bonus. Moreover, as for the second mover the bonus still is unattractive under 1st-order beliefs (since the opponent is assumed to best-respond to an L0 player and hence

---

[4]For an insured player termination at any (if first mover) or any non-terminal node (if second mover) is strictly dominated.

always plays pass), we would also expect fewer bonus offers to be made by her. Both predictions of our behavioural model are consistent with the data.

Summing up, in our view the main contribution of the paper is twofold. First of all, the experimental data provide further support for the view that (sequential) mechanisms designed to implement a socially desirable outcome may fail in practice if they implement this outcome only under rational or higher-order beliefs as in case of the insurance in the centipede game. In this respect, we believe that our experiment delivers an important insight for applied mechanism design. Secondly, apart from these more general aspects, the more detailed analysis in terms of higher-order beliefs not only provides some clues as to the subjects average level of iterated backward reasoning. It also demonstrates that level-k-thinking models, if suitably adapted, may contribute to the understanding of individual behaviour in extensive form games, at least if the respective games are comparably simple.

The rest of the paper is structured as follows. In Section 2, we describe and analyse the different treatments of our experiment. In Section 3, we report the data and discuss the results of the experiment. Section 4 concludes.

2 Experimental Design

To facilitate the exposition, this section is divided into two parts. In the first part, 2.1, we describe the general experimental set-up as well as the two instruments to be tested and we derive the *rational*[5] prediction for these cases. In the second part, 2.2, we specify the remaining details of the different experimental treatments.

2.1 General Description and Theoretical Analysis

For our experiment we consider three different versions of the centipede game: the standard centipede game (henceforth CP-game), a CP-game with insurance and a CP-game with a bonus option. In the following we explain these treatments in detail.

---

[5]Here and in the following the *rational* prediction is understood to be the subgame perfect Nash equilibrium.

2.1.1 The Standard CP-Game

As a basis for our study, we use the 8-move centipede game depicted in Figure 1, denoted by G_S. Different from earlier studies (e.g. McKelvey and Palfrey, 1992; Nagel and Tang, 1998), we use linearly instead of exponentially increasing payoffs. This choice appears to be better suited for the interpretation of the game in terms of a joint venture as indicated in the introduction. Also, under linearly increasing payoffs crazy type explanations for passing (Kreps et al., 1982) are less plausible.[6] The *8-move* specification is chosen so as to leave room for gains which are costly to achieve — given the threat of early termination — while at the same time avoiding "habitual passing" in the beginning which may occur if the game is too long.
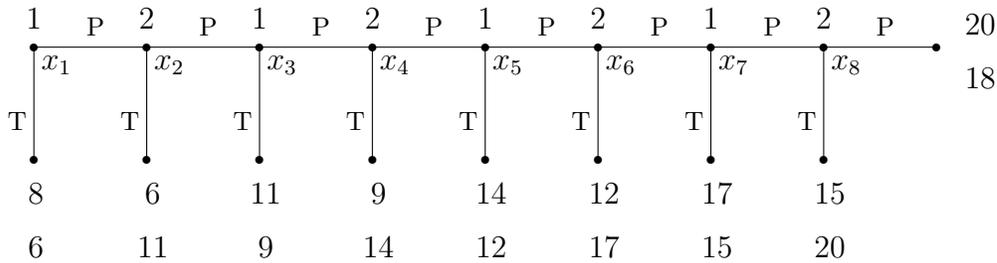

Figure 1: G_S. Player 1 (2) moves at nodes x_k (x_{k+1}) for k odd and can choose the actions T (*terminate*) or P (*pass*). At each terminal node player 1 (2) obtains the upper (lower) payoff.

The rational prediction for this standard CP-game, G_S, is straightforward. We state it without proof. The equilibrium concept employed for this as well as for the other cases is that of a subgame perfect Nash equilibrium, henceforth SPNE.[7]

Rational Prediction (RPS):

In the unique SPNE of G_S each player chooses *terminate* (T) at all her decision nodes. Hence, the game is terminated at node x_1.

---

[6]More altruists or crazy types would be needed as the potential gains from matches with an altruist are lower.

[7]Incidentally, in G_S the Nash equilibrium outcome is unique and coincides with the subgame perfect Nash equilibrium outcome.

2.1.2 The CP-Game with Insurance

In this version the CP-game is extended to a two-stage game. In stage 1 both players simultaneously decide whether to purchase an insurance at price c (the exact value of c is specified later). Once both players have decided whether to buy an insurance or not, they are informed about the decision of the opponent and the game moves to stage 2, where the CP-game is played. The final payoff for a player is the payoff obtained in the CP-game minus c, if the player is insured, plus an indemnity $2 + \varepsilon$ with $\varepsilon > 0$, if the player is insured and the CP-game is terminated by the opponent. The resulting game is denoted by G_I. Figure 2 illustrates the final payoffs for the case where only player 2 is insured. Observe that the insurance never makes a loss if $c \geq 2 + \varepsilon$.
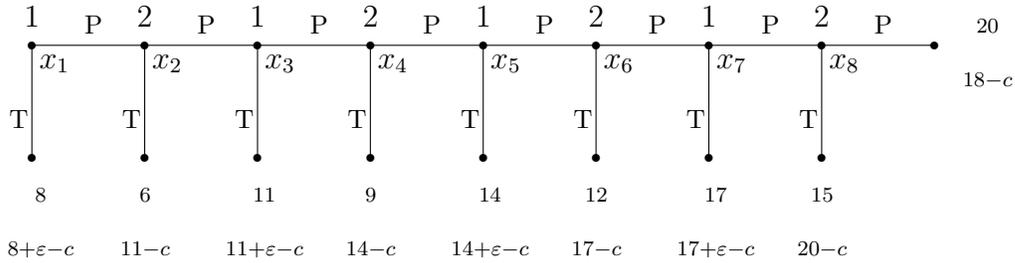

Figure 2: Final payoffs in G_I for the case where only player 2 is insured.

The potential benefit of the insurance lies in the fact that, once a player is insured, *pass* becomes the dominant action for this player at all nodes except node x_8, where it is player 2's move and 2 optimally chooses *terminate* even if she is insured. As the players' insurance decisions are revealed prior to the play of the CP-game, purchasing the insurance can be viewed as a credible signal that the player will pass at all nodes (except x_8). The complete rational prediction for G_I is the following (its derivation can be found in the appendix):

Rational Prediction (RPI):
Insurance Decision: If $c \in (0, \varepsilon)$, then there is a unique SPNE, in which player 1 is insured with probability 1 and player 2 is uninsured. If $c \in [\varepsilon, 9 + \varepsilon]$, then in any SPNE of G_I, either one of the two players is insured with probability 1

and the other player is not, or both players mix and the insurance is bought with probability

$$p_1^* = 1 - \frac{c}{11 + \varepsilon} \quad \text{and} \quad p_2^* = 1 - \frac{c - \varepsilon}{9}$$

for player 1 and player 2, respectively. If $c \in (9 + \varepsilon, 11 + \varepsilon)$, the unique SPNE is given by only player 2 buying the insurance. If $c > 11 + \varepsilon$ there is a unique SPNE in which no player is insured. Finally, in the knife-edge case where $c = 11 + \varepsilon$, in any SPNE either no player is insured or only player 2 is insured.

CP-Game: The SPNE-outcomes of the CP-game at stage 2, given the players' insurance decision, are summarised in Table 1.

Insured	Strategies	Payoff for 1	Payoff for 2
Both	T at x_8	$17 + \varepsilon - c$	$20 - c$
Only Player 1	T at x_8	$17 + \varepsilon - c$	20
Only Player 2	T at x_7	17	$17 + \varepsilon - c$
None	T at x_1	8	6

Table 1: SPNE-outcomes of the CP-game given the players' insurance decision. Here, "T at x_k" is to be understood as "P at all nodes x_m, $m < k$, and T at all nodes x_m, $m \geq k$."

2.1.3   The CP-Game with Bonus Option

In the third version of the CP-game, the players' insurance decision is replaced by the option to offer the opponent the transfer of a bonus $b > 2$ which is paid only in case the opponent does not terminate the game. The size of the bonus is exogenously given. The offer again has to be decided upon simultaneously and prior to the play of the CP-game. The bonus offer decisions are observed by both players before the game moves to stage 2, where the centipede game is being played. Any offer is binding. The resulting game is denoted by G_B; see Figure 3 for an illustration of the final payoffs in the case where only player 1 has offered a bonus to player 2.

The potential benefit of the bonus option is straightforward. For $b > 2$, passing at all nodes is a dominant strategy for both players whenever at least one
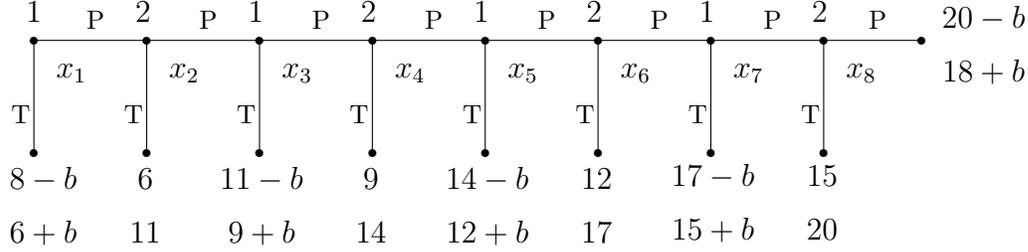
Figure 3: Final payoffs in G_B for the case where only player 1 has offered to pay the bonus b if player 2 does not terminate the game.

player has offered the bonus. The complete rational prediction for this version of the CP-game is given below (see the appendix for a derivation):

Rational Prediction (RPB):
Bonus Decision: If $b \in (2,5)$, the unique SPNE is given by only player 1 offering the bonus. If $b \in [5,12]$, there are three SPNE: either only player 1 (player 2) offers the bonus with certainty or both players mix in their bonus decision with

$$p_1^* = 1 - \frac{b}{14} \quad \text{and} \quad p_2^* = \frac{12 - b}{7}.$$

If $b \in (12,14)$, the unique SPNE is given by only player 2 offering the bonus. If $b > 14$ no player offering the bonus is the unique SPNE. Finally, in the knife-edge case where $b = 14$, in any SPNE either only 2 offers the bonus or no player offers the bonus.
CP-Game: The SPNE-outcomes of the CP-game at stage 2, i.e. given the players' bonus offers, are summarised in Table 2.

2.2 Specifications for the Experiment

For the experiment, we implemented three treatments: a standard treatment (an implementation of G_S); an insurance treatment (an implementation of G_I with $c = 2.5$ and $\varepsilon = 0.5$); and a bonus treatment (an implementation of G_B with $b = 2.5$). The parameters are chosen such that a player who is insured has the same incentive to play pass as a player who was offered a bonus: in both cases the difference in payoffs between immediate termination and termination by the

b is offered by	Strategies	Payoff for 1	Payoff for 2
Both	P at x_8	20	18
Only Player 1	P at x_8	$20 - b$	$18 + b$
Only Player 2	T at x_8	$15 + b$	$20 - b$
None	T at x_1	8	6

Table 2: SPNE-outcomes of the CP-game given the players' bonus offers. Here, "P at x_8" is to be understood as P at all nodes. As before, "T at x_k" is to be understood as "P at all nodes x_m, $m < k$, and T at all nodes x_m, $m \geq k$."

opponent at the next node is 0.5. See Table 3 for a summary of the specifications and the respective rational predictions.

The computerized experiment was conducted at the Institute for Empirical Research in Economics at the University of Zurich.[8] Subjects for the experiment were recruited among students from the University of Zurich and the Swiss Federal Institute of Technology Zurich (ETHZ).[9] The data for each treatment were obtained in two separate sessions. In each session there were 20 subjects. Subjects in all treatments were equipped with detailed instructions for the experiment and had to answer several test-questions before the experiment was started.[10] In order to make sure that all subjects had understood the rules of the game we did not start the experiment until all subjects had correctly answered all questions. All subjects in a session, then, were randomly assigned to one of two equally large groups: the *red* group, which had the role of player 1, and the *blue* group, which had the role of player 2. It was made clear to the subjects that they would keep their respective label, red or blue, throughout the entire experiment and that each member of the red group would play the respective CP-game version exactly once against each member of the blue group. Hence, in each session the game was repeated ten times, referred to as periods, and no subject ever played the game against a previous opponent or against the opponent of a previous op-

---

[8]The experiment was programmed and conducted with the experimental software z-tree (Fischbacher, 1999).

[9]Subjects were recruited with the online recruitment system ORSEE (Greiner, 2004).

[10]All supplementary material (instructions, test-questions, questionnaire, etc.) is available from the authors on request.

Treatment	Parameters	Rational Prediction and Payoffs
Standard	—	Immediate termination at x_1 with payoffs $= (8, 6)$
Insurance	$\varepsilon = 0.5$ $c = 2.5$	Three Nash equilibria: Only player 1 insured with payoffs $(15, 20)$. Only player 2 insured with payoffs $(17, 15)$. Both players mix with payoffs $(15, 16.96)$. Termination according to who is insured: Only player 1 or both $\Rightarrow x_8$ Only player 2 $\Rightarrow x_7$ No player $\Rightarrow x_1$
Bonus	$b = 2.5$	Only player 1 offers the bonus. Permanent passing with payoffs $(17.5, 20.5)$.

Table 3: Overview over the specification for the different treatments and the respective rational predictions.

ponent. Actual interaction was anonymous. After each interaction, subjects were informed about the results of all their previous interactions.

After the computer based part of the experiment, subjects were asked to answer a questionnaire about their behaviour in the experiment as well as about some general social and risk attitudes. The questionnaire data are reported in the sequel only insofar as we consider them pertinent to the understanding of the subjects' behaviour in the experiment. At the end of each experimental session, all experimental payoff units earned by the subjects were converted into Swiss Francs at a rate of 10:1 (i.e. 10 experimental payoff units were exchanged for CHF 1).[11] Table 4 summarizes the duration and average earnings of the subjects in each treatment.

---

[11]CHF 1 is approximately equal to €0.64 or $0.81.

Treatment	Duration (incl./excl. Instructions)	Average Payoff (CHF)
Standard	25 Min./8 Min.	15.60
Insurance	55 Min./13 Min.	15.58
Bonus	55 Min./23 Min.	17.60

Table 4: Overview over the duration and average earnings of the subjects in the different treatments.

3   Results

In the following, we report and discuss the data obtained in each of the three treatments specified above. Each treatment was conducted with 40 subjects (20 subjects per session, 2 sessions per treatment) resulting in 200 observations per treatment.

3.1   The Standard Treatment

First, we report the results from the standard treatment which will serve as the benchmark case for our study.

Data
Figure 4 shows the relative frequencies with which the game terminated at the different nodes, aggregated over all periods. Average CP-game payoffs for subjects assigned to the role of player 1 (first-mover) were CHF 14.95; average CP-game payoffs for subjects assigned to the role of player 2 were CHF 16.24.

Different from McKelvey and Palfrey (1992), we do not find any evidence for learning towards the rational solution over time. The cumulative distribution of termination in later periods does not dominate the distribution in earlier periods in the sense of first order stochastic dominance (cf. Figure 5). In fact, median termination remains quite stable between node x_6 and x_7. Hence, we can safely aggregate the data over all periods.[12]

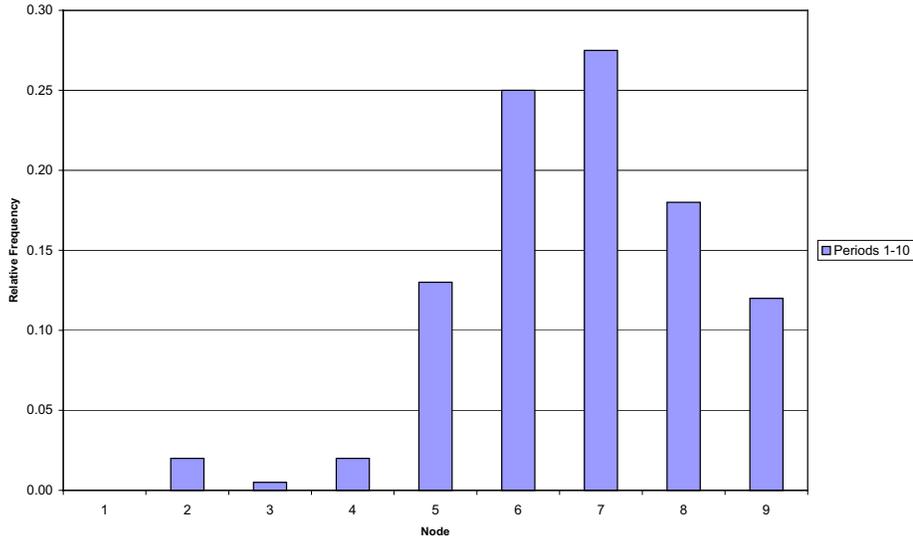[12]Also individually, small variations are rather indicative of experimentation than of (strategic) learning.

Figure 4: Relative termination frequencies for the standard treatment aggregated over all periods. Termination at node 9, i.e. x_9, is understood to mean that both players always played *pass*.
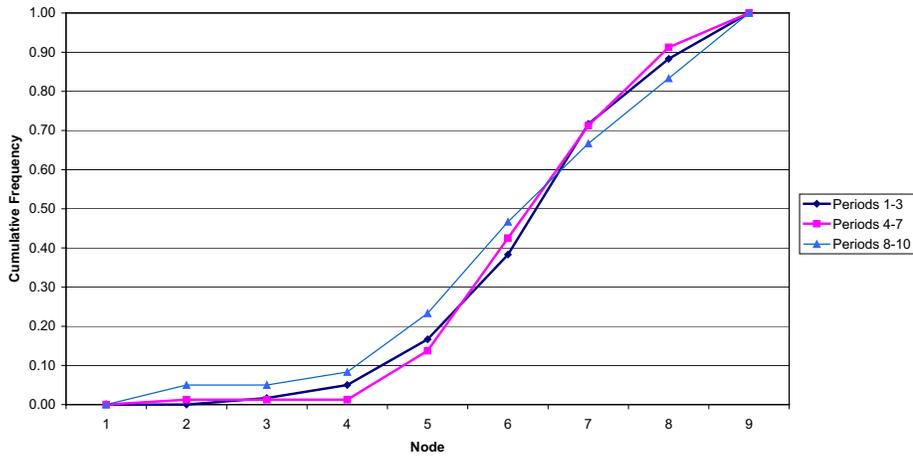

Figure 5: Cumulative termination frequencies for the standard treatment aggregated over periods 1–3, 4–7, 8–10.

Discussion

The observed termination frequencies are consistent with a simple behavioural model in terms of level-k thinking (cf. Stahl and Wilson, 1994, 1995; Nagel, 1995; see also Ho et. al, 1998; Costa-Gomes et al., 2001; Crawford and Iriberri, 2007a, 2007b). However, as these models are commonly used for the analysis of normal form games, we need some adjustments to the specific sequential structure

of the CP-game. In particular, we assume that non-strategic players (level-0), L0 for short, do not randomise over all available strategies, which is the most common assumption, but instead always play pass. This all-pass-assumption for non-strategic L0 players seems plausible as, given the similarity of all pass-take decisions, playing pass is arguably *the best one can do* without thinking strategically as termination is costly in terms of foregone payoffs.[13] All other players, then, are assumed to best-respond to finite-order beliefs, where a belief is classified as being of order k if the respective player expects her opponent to play a best-reply against a player with a belief of order $k - 1$. This latter specification again slightly differs from the literature on level-k thinking in normal form games where it is commonly assumed that level-k players simply best respond to level-$(k - 1)$ players (e.g. Nagel, 1995). Under the common definition, however, it would be impossible to distinguish between non-strategic play and a best-response to a 1st-order belief in player 1's behaviour at x_7 (in both cases player 1 always plays pass). Consequently, for player 2 one could not distinguish between best-response play to a 1st or 2nd-order belief at x_6 and so on. Therefore, we prefer our type classification which gives a one-to-one relation between types and termination behavior; i.e. it ties the level of thinking to the steps of iterated backward reasoning.

With the above type classification, player 1 with 1st-order beliefs expects player 2 to pass until the last node and terminate there. Player 1's best-reply, then, is to pass until the next but last node, i.e. x_7, where she terminates the game. Similarly, with 1st-order beliefs, player 2 expects player 1 to always pass and, hence, does not terminate before the last node. By the same token, under 2nd-order beliefs player 2 terminates at x_6 since she expects player 1 to terminate at x_7. Analogously, player 1 terminates at x_5 under 2nd-order beliefs since he expects player 2 to terminate at x_6; etc. Accordingly, for the 8-move centipede game considered here, 4th-order beliefs by player 1 already lead to fully rational behaviour, i.e. immediate termination of the game; see Table 5 for a summary.

From Table 5 we can derive predictions for the termination of the game under different beliefs of the players. For example, if player 1 has 1st-order beliefs

---

[13]This, of course, presumes that both a more detailed differentiation of a player's own decision nodes and considerations about the opponent's behaviour are signs of "strategic" thinking.

Type	Red	Blue
L0	all pass	all pass
1st-order beliefs	x_7	x_8
2nd-order beliefs	x_5	x_6
3rd-order beliefs	x_3	x_4
4th-order beliefs	x_1	x_2

Table 5: Intended termination nodes according to the order of beliefs for player 1 (red) and player 2 (blue).

and player 2 is either non-strategic, $k = 0$, or has 1st-order beliefs, we predict the game to terminate at x_7, which corresponds to the peak of the observed distribution (cf. Figure 4).

We use aggregate termination frequencies in order to assess the percentage of observed behaviour which is consistent with 1st-,..., 4th-order beliefs or with non-strategic permanent passing. The resulting type-distribution we then take as a proxy for the degree of rationality of the subjects in the respective roles (red / blue) also for later treatments. Of course, in doing so we neglect variations in individual behaviour over time. However, as we do not find any evidence for learning between periods (in any of the treatments), we attribute these changes to occasional experimentation and not to increased rationality. As we will see, behaviour in later treatments is supportive of this assumption.

To actually determine the average type distribution, we have to deal with those subjects for whom intended play remains ambiguous because of termination by their opponent. We assume here that the behaviour of these subjects at unreached nodes on average would have been the same as observed play at these nodes.[14] The type distribution thus derived is shown in Table 6.

In total, we find that aggregate termination frequencies in the standard game are consistent with 85.95% of the red and 66.90% of the blue players being non-strategic, $k = 0$, or having 1st-order beliefs. This observation is also consistent

---

[14]As an example consider the case of the blue player. Of those blue players who reach x_8, 60% terminate the game there. Thus, we assume that of those who pass at x_6, 60% would have terminated the game at x_8 although not all get that far due to termination by the red player at x_7.

Type	Red	Blue
L0	44.84%	26.76%
1st-order beliefs	41.11%	40.14%
2nd-order beliefs	13.54%	29.09%
3rd-order beliefs	0.51%	2.01%
4th-order beliefs	0%	2%

Table 6: Extrapolated type distribution for player 1 (red) and player 2 (blue).

with the individual data, according to which only 3 out of 20 (8 out of 20) subjects in the role of the red (blue) player exhibited a termination behaviour that necessitates more than 1st-order beliefs in more than 3 instances to be consistent with our model.[15]

It may seem puzzling that according to Table 6 blue players on average appear to be more rational, although all participants are randomly drawn from the same subject pool.[16] In fact, there are more red than blue L0 players and more blue than red players with beliefs of order $k > 1$; the percentage of red and blue players having 1st-order beliefs is about the same. Yet, there are two potential rationales that can account for this particularity in the type distribution. First of all, due to our choice of type classification, it is likely that we have underestimated the degree of rationality among red players. In particular, since observed behaviour does not allow us to distinguish between L0 first-movers and those who best respond to an L0 second-mover, we assigned both types to the class of L0 players. However, it may well be that some of the (so classified) L0 players in fact play pass as a best-reply to an L0 opponent. As a similar argument also applies to higher-order beliefs, the estimated type-distribution for red players as a whole is likely to be slightly biased towards lower degrees of rationality if we compare it to the type-distribution for blue players. Secondly, being less rational in terms of the order of beliefs used, i.e. exerting less effort in analysing the game, is always more costly for player 2 (blue) than for player 1 (red). To see this, consider the

---

[15]Few instances could be attributed to simple experimentation. We chose 3 as a threshold, but other choices (2 or 4) would not change the numbers in any significant way.

[16]This effect can also be found in the data of McKelvey and Palfrey (1992).

case of non-strategic passing. An L0 red player, passing at x_7, still can, and occasionally will, encounter a blue player passing at x_8, i.e. permanent passing by the red player is a best-reply to an L0 opponent. By contrast, permanent passing is always a dominated strategy for blue, even if blue plays against an L0 red. Hence, the structure of the centipede game appears to be such that it puts more pressure on blue players to exhibit a comparably high degree of rationality in their termination behaviour. On further reflection, the observed differences in the type distribution, thus, are less surprising. We do not pursue this issue any further, though, as only the type's termination behaviour is important for the ensuing discussion.

3.2   The Insurance Treatment

Next, we report and discuss the results from the insurance treatment.

Data

In 147 out of 200 pairs none of the two players purchased the insurance (cf. Figure 6). Aggregated over all periods, red players insured themselves in 15 out of 200, blue players in 41 out of 200 cases. In three instances this resulted in both players being insured. Moreover, only 4 out of 20 red players and 10 out of 20 blue players used the insurance at least once.[17] Thus, the insurance was completely dismissed by a majority of 26 out of 40 subjects.

Accordingly, aggregate relative termination frequencies (cf. Figure 7) are almost identical to the standard treatment; using a Kolmogorov-Smirnov two-sample test, we found no significant difference. Also, there is again no clear evidence of learning towards earlier termination.[18]

---

[17]One of the red players used it in all ten periods though.

[18]As the learning-towards-Nash aspect is not the focal point of our analysis, we do not report the respective data in further detail.
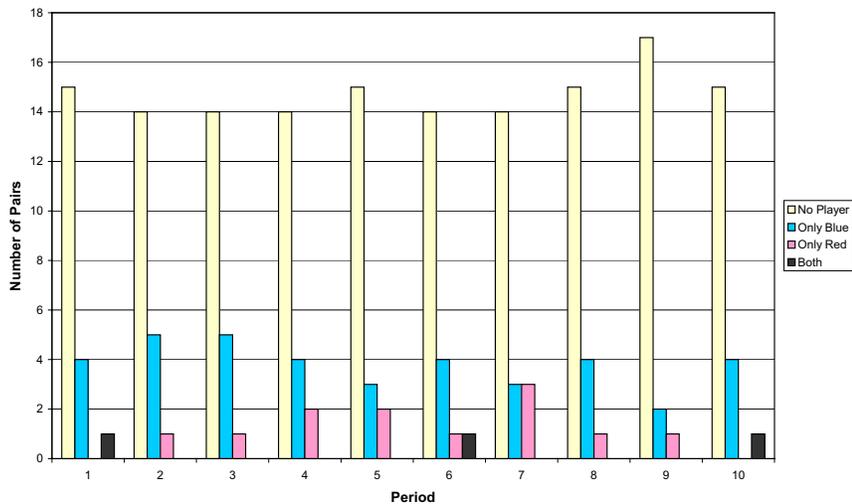
Figure 6: Number of pairs per period where no player, only player 2 (blue), only player 1 (red), or both players were insured.
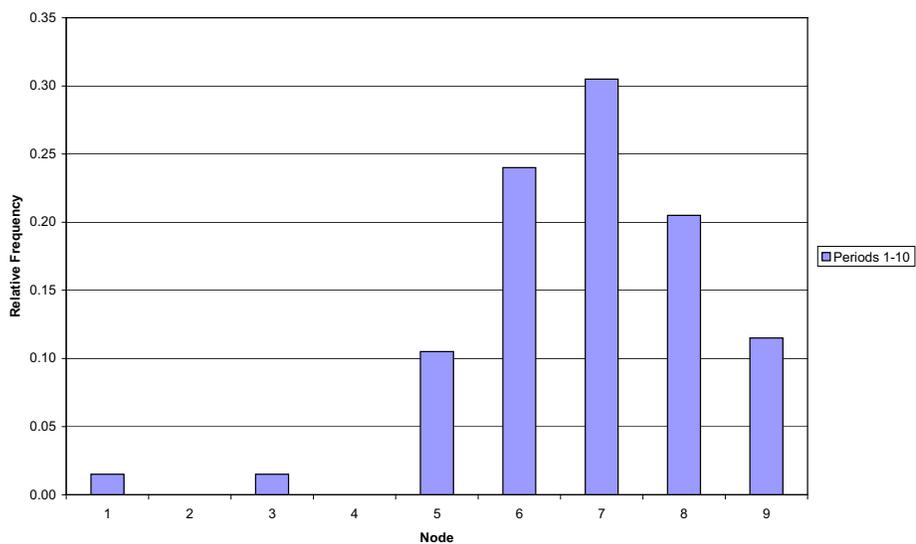

Figure 7: Relative termination frequencies for the insurance treatment aggregated over all periods.

Interestingly, despite its infrequent usage, the insurance option did not completely fail its purpose. In fact, conditional on only one player being insured, termination has a clear peak at the rational prediction, i.e. at x_8 if only player 1 is insured and at x_7 if only player 2 is insured (cf. Figure 8). If no player is insured, termination is centered around nodes x_6 and x_7, similar to the standard treatment.
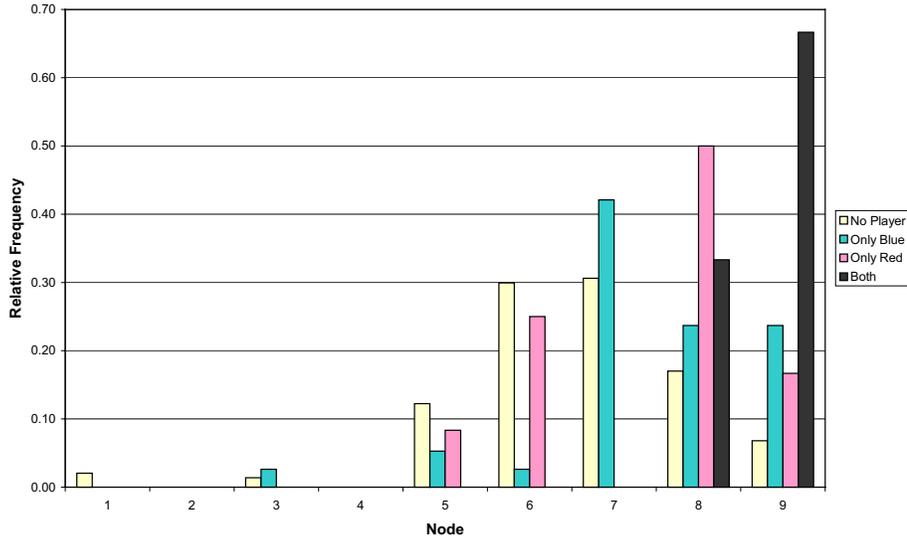

Figure 8: Relative termination frequencies for the insurance treatment conditional on no player, only player 2 (blue), only player 1 (red), or both players being insured, aggregated over all periods 1–10.

Yet, due to the occasional too early termination of players (even insured ones), introducing the insurance option does not pay in terms of payoffs (cf. Table 7).

Player	Average payoff (CHF)		
	if insured	if not insured	unconditional
1/Red	14.83	15.16	15.13
2/Blue	15.44	16.19	16.04

Table 7: Average payoffs earned by the subjects in the insurance treatment.

Discussion

Two results are particularly noteworthy: first, that the insurance option is used so little; and second, that in most cases it is player 2 (blue) who buys the insurance. In fact, at least one player is insured for 53 out of 200 pairs only. Moreover, in 41 out of these 53 cases player 2 (blue) is insured and only in 15 cases player 1 (red) is insured.[19] Also, out of the 14 subjects who insured themselves at least once, 10 had the role of player 2 (blue) and only 4 had the role of player 1 (red); 5 blue players and 2 red players used the insurance more than once. These observations are in clear contrast to either of the rational predictions for this game; even in the mixed equilibrium, both players buy the insurance with a probability larger than 80%.

However, these observations are to a large extent consistent with the belief-based model introduced in the previous subsection. To see why, assume that a player buys the insurance only if the expected delay in game termination is sufficiently large so that she expects a higher payoff compared to the situation where none of the players is insured. Obviously, in this case the insurance is worthless for a player with 1st-order beliefs since the opponent is believed to play a best-reply to a L0 player who always plays pass. Hence, under 1st-order beliefs the opponent's behaviour is independent of the own insurance status so that the insurance is expected to entail a loss rather than a gain. A player therefore needs at least 2nd-order beliefs in order to realise the potential value of the insurance. Under 2nd- or higher-order beliefs the insurance is expected to signal the opponent a change in behavior compared to the case of no insurance, namely that the insured player always plays pass (if red) or plays pass until the last decision node and terminate there (if blue). If the opponent best-responds to this dominant strategy of an insured player, game termination is delayed.

More precisely, an insured red player with at least 2nd-order beliefs expects her blue opponent to terminate the game at x_8. So red obtains a higher payoff with the insurance than without it; in the latter case she expects the game not to go beyond x_5 (cf. Table 8). A blue insured player with at least 2nd-order beliefs expects her red opponent to terminate at x_7. Without insurance she expects the game to terminate at x_6 under 2nd-order beliefs, at x_4 under 3rd-order beliefs, or

---

[19]Recall that there are three pairs in which both players are insured.

at x_2 under 4th-order beliefs (cf. Table 8). Hence, only under 3rd- or 4th-order beliefs the blue player considers the delay in game termination to be sufficiently large in order to buy the insurance. As regards the non-strategic L0 players, we assume that they do not buy the insurance as, strategic thinking aside, it offers no benefit but is associated with a cost (cf. Table 8).

Type	Red		Blue	
	w/o insurance	w. insurance	w/o insurance	w. insurance
L0[20]	all pass (**20**)	all pass (17.5)	all pass (**18**)	all pass (15.5)
1st-order	x_7 (**17**)	x_8 (15)	x_8 (**20**)	x_8 (17.5)
2nd-order	x_5 (14)	x_8 (**15**)	x_6 (**17**)	x_7 (15)
3rd-order	x_3 (11)	x_8 (**15**)	x_4 (14)	x_7 (**15**)
4th-order	x_1 (8)	x_8 (**15**)	x_2 (11)	x_7 (**15**)

Table 8: Expected termination node without (w/o) and with (w.) *own* insurance and corresponding own payoffs (including insurance costs and benefits, if applicable) for red and blue dependent on the order of beliefs.[20] Payoffs for the respective preferable option are in bold type.

Taking the type distribution obtained in the standard treatment as the benchmark (cf. Table 6), we thus expect red (blue) players to use the insurance in about 14% (4%) of all instances. According to our data, red players use the insurance in 7.5% of all cases which is below our prediction. Individually, however, 20% of the red players use the insurance at least once and 10% use it more than once. Although the post experimental questionnaire gives few hints at framing effects (Andreoni, 1995), they nevertheless offer a potential rationale for the comparably small discrepancy. In particular, the insurance decision, arguably, may have been prone to *negative* framing effects, especially for the first-moving red; purchasing insurance after all may be construed as a lack of confidence in the opponent's good will and may thus obliterate any positive appeal to reciprocity due to initial

---

[20]As regards the non-strategic L0 players, specifying expectations, of course, is awkward as any expectation about termination implies some form of strategic thinking. Accordingly, the "all-pass" expectation specified in the table should not be understood as a "conscious expectation" but rather as the logical equivalent to the assumed behaviour of L0 in the CP-game.

passing. All in all, we think that the observations for red are fairly consistent with the prediction from the model.

As concerns observed usage frequencies for blue players, these are notably higher than predicted. In aggregate blue players use the insurance in 20.5% of all cases and 25% of the blue players use the insurance more than once (prediction $\approx 4\%$; cf. Table 6). However, the high usage frequencies for blue players can be explained by two possible mistakes on the part of a blue player with 2nd- or higher order beliefs. First, a blue player may simply compare the expected game termination with and without insurance and buy the insurance whenever it delays the termination of the game. In this case she buys the insurance already under 2nd-order beliefs although her payoff is lower than without insurance due to a delay by one node only (x_7 instead of x_6, cf. Table 8). Second, blue may falsely believe that the red opponent always plays pass if blue is insured. This error may arise as follows: When contemplating to buy the insurance a blue player with 2nd- or higher order beliefs may argue that the insurance signals the opponent that he will always play pass, while in fact he will not play pass at the last node. The red opponent's best-reply then is to always play pass which means the game terminates at x_8. Hence, under both mistakes the insurance is already valuable for blue players with 2nd-order beliefs. Including this type in our prediction, we would expect up to 33% of the blue players to buy the insurance which is indeed close to what we observe in the data.[21]

Besides our behavioral model, there is an alternative, more rational explanation for the fact that most subjects do not purchase the insurance. As can be seen from Table 7, average payoffs were higher for subjects if they were not insured. Hence, if a subject either has a correct belief about the payoffs to be expected with and without insurance or if he or she has tested the insurance in some periods and payoffs were observed to be too low, then it is rational not to buy the insurance. Also, given theoretical predictions (cf. Table 3), neither player could increase her expected payoff beyond the empirical mean in the standard treat-

---

[21]Note that negative framing effects, as cited above, should only play a negligible role in case of blue. Not only is the evidence for such effects from the post experimental questionnaire feeble. If present at all, we also would expect the effect to be less negative for blue as second-movers are more likely to feel "in need" of an insurance (or may expect to be perceived as such), given that first-movers eventually can always terminate the game in their favour.

ment by taking the insurance herself.[22] And, indeed, only 5 out of 14 subjects who insure themselves at least once buy the insurance more than 4 times. In view of the subjects' responses to the post experimental questionnaire, however, the rational explanation for the failure of the insurance is questionable. If subjects had rejected the insurance option due to optimistic expectations about expected payoffs in the standard game, they could have reported so in their answers to the questionnaire. Yet, only one subject comments on the own behaviour in a way which indicates that the insurance was dismissed due to roughly correct expectations about the outcome without insurance. The majority of subjects instead indicate that they did not realize that the insurance may delay the termination of the game which is in line with our behavioural model.

3.3   The Bonus Treatment

Finally, we report the results from the bonus treatment. Recall that in the bonus treatment, prior to the play of the CP-game subjects could offer their opponent a bonus $b = 2.5$ for not terminating the game.

Data

Different from the insurance, the opportunity to offer a bonus to the opponent for not terminating the game is readily seized (cf. Figure 9). Player 1 (red) offered the bonus in 133 out of 200 instances, with 19 out of 20 subjects offering the bonus at least once; 15 out of 20 subjects offered it strictly more than 3 times. Player 2 (blue) offered the bonus in 66 of 200 instances, with 15 out of 20 subjects offering the bonus at least once; 9 out of 20 subjects offered the bonus strictly more than 3 times. Together this resulted in a total of 154 of 200 pairs where at least one player had offered the bonus. As we can see from Figure 9, in all periods except the first only player 1 (red) offered the bonus for the majority of pairs, which is in accordance with the theoretical prediction.

Moreover, in the vast majority of cases, if a bonus offer is made, the CP-game is terminated according to the rational prediction or even later, i.e. no earlier than node x_8 (cf. Figure 10). Only if no bonus is offered, relative termination frequencies are centered around nodes x_6 and x_7 as in the standard treatment.

---

[22]We treat the average payoff of 14.95 for the red player in the standard case as equal to 15.
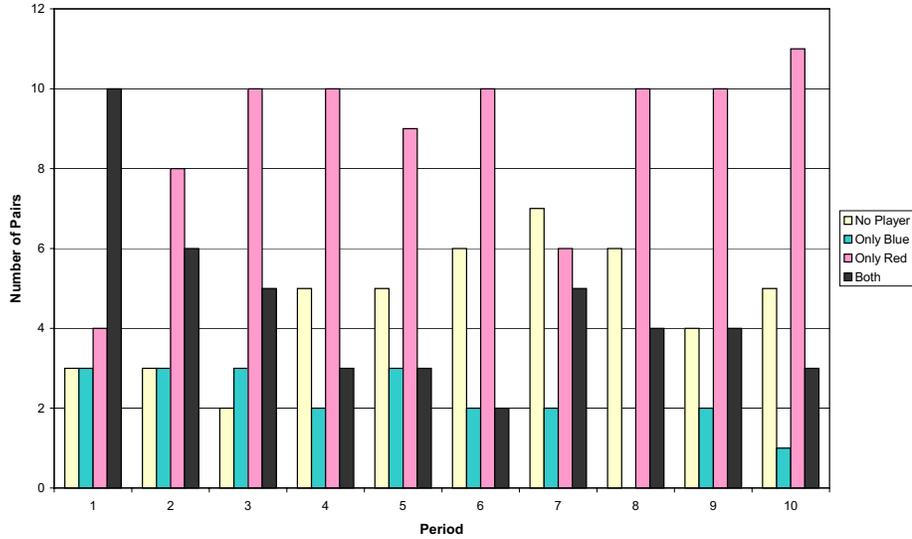
Figure 9: Number of pairs in the bonus treatment, where no player, only player 2 (blue), only player 1 (red), or both players offered a bonus payment.
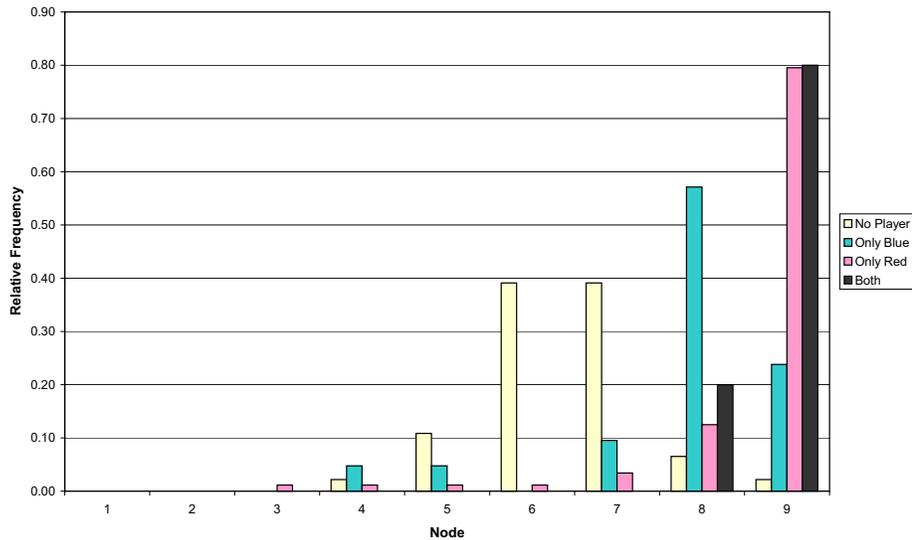

Figure 10: Relative termination frequencies for the bonus treatment conditional on no player, only player 2 (blue), only player 1 (red), or both players offering a bonus, aggregated over all periods 1–10.

Accordingly, as the bonus offer is used so frequently, the CP-game is terminated significantly later in the bonus treatment than in the standard or insurance treatment (p-value $\ll 0.01$ in a one-tailed Kolmogorov-Smirnov test; cf. Figure 11).
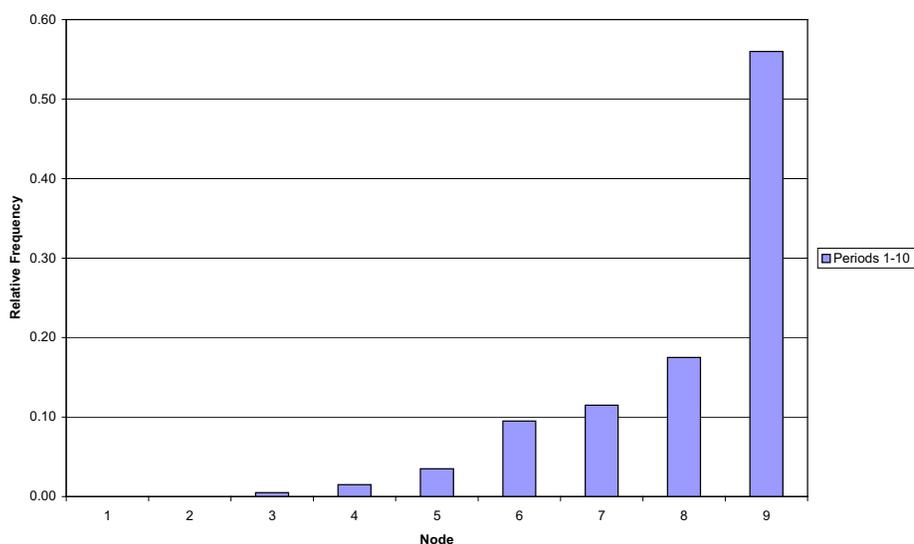

Figure 11: Relative termination frequencies for the bonus treatment aggregated over all periods.

Finally, for red players payoffs were considerably higher if they offered a bonus than if this option was dismissed (cf. Table 9).

Player	Average payoff (CHF)		
	if subject offered bonus	if subject offered no bonus	unconditional
1/Red	17.68	15.64	17.00
2/Blue	17.37	18.60	18.20

Table 9: Average payoffs earned by the subjects in the bonus treatment.

Discussion

Our data show that the bonus option is effective in that the bonus is offered by at least one of the players in 77% of all pairs; in 66.5% (33%) of all possible cases the red (blue) player offered the bonus. Moreover, if the bonus is offered, the number of terminations prior to x_8, the terminal node, is negligible (11 of of 154 cases, i.e. 7.1%).

These observations are consistent with the previous specification of our belief based behavioural model. Again, assume that a player offers the bonus if the expected delay in game termination is sufficiently large such that she expects a higher payoff compared to the situation where no player offers a bonus. In this case, all red players with at least 1st-order beliefs will note that offering the bonus renders permanent passing the dominant strategy for blue (cf. Table 10). Without the bonus, all strategic red players ($k > 0$; 55% according to our estimation in Table 6) expect blue to terminate at x_8 – at the latest – so that it is profitable to offer the bonus. For a blue player, by contrast, the bonus is valuable only under 2nd- or higher-order beliefs (cf. Table 10). With 1st-order beliefs blue expects red to always play pass even without being offered a bonus. Thus, only blue players with 2nd- or higher-order beliefs will offer the bonus (33% according to our estimation in Table 6). Finally, non-strategic L0 players will not consider offering a bonus to their opponent (cf. Table 10) as it involves a cost but provides no benefit to them. To see this observe that in order to recognize the benefit of the bonus, a player has to analyse the opponents' termination behaviour, i.e. she has to think strategically which is ruled out for L0 players by definition. These predictions are indeed close to the observed frequencies for bonus offers given by 66.5% for red (we predicted 55%) and 33% for blue which is, incidentally, identical to our prediction. This result also holds true if take into account potentially positive framing effects in conjunction with the bonus offer (cf. Andreoni, 1995). In this case, predicted figures for red and blue would have to be slightly increased; yet again we would expect the increase to be moderate as post experimental questionnaires give no conclusive hint at the presence of framing effects.

Type	Red		Blue	
	w/o bonus	w. bonus	w/o bonus	w. bonus
L0[23]	all pass (**20**)	all pass (17.5)	all pass (**18**)	all pass (15.5)
1st-order	x_7 (17)	all pass (**17.5**)	x_8 (**20**)	x_8 (17.5)
2nd-order	x_5 (14)	all pass (**17.5**)	x_6 (17)	x_8 (**17.5**)
3rd-order	x_3 (11)	all pass (**17.5**)	x_4 (14)	x_8 (**17.5**)
4th-order	x_1 (8)	all pass (**17.5**)	x_2 (11)	x_8 (**17.5**)

Table 10: Expected termination node without (w/o) and with (w.) *own* bonus offer and corresponding own payoffs (including bonus cost, if applicable) for red and blue dependent on the order of beliefs.[23] Payoffs for the respective preferable option are in bold type.

4 Conclusion

In this paper, we have proposed two modifications of the centipede game aimed to achieve higher aggregate welfare, i.e. later termination. These modifications were given by an insurance which slightly overcompensates losses due to termination by the opponent, and by an option to offer a bonus to the opponent for not terminating the game. Both modifications were designed so as to theoretically induce passing until (close to) the end once they were utilised. Different from the bonus option, which is more direct, the insurance banks on a signalling effect as buying the insurance indicates that the insured player intends to always pass.

Empirically, it turned out that subjects fail to grasp the signalling aspect of being insured and essentially ignore the insurance option. By contrast, the bonus, which can readily be interpreted as a price for permanent passing by the opponent, was heavily used and led to significantly later termination of the CP-game as compared to the standard version. Thus, the results show that it is considerably easier for subjects to grasp effects based on strict dominance (bonus) than effects based on iterated reasoning (insurance).

But more can be said. As we have argued, the results can be rationalised by a

---

[23]Cf. footnote 20.

simple behavioural model based on the assumption that players essentially best-respond to finite-order beliefs. In particular, the data obtained in the standard treatment suggest that the majority of subjects choose their strategy based on beliefs of order at most 1. With such beliefs, however, subjects cannot realise the signalling value of the insurance, which requires at least two steps of iterated reasoning for the first and 3 steps of iterated reasoning for the second mover. Yet, they can grasp the more immediate effect of the bonus. Both the empirical data and the answers from the questionnaire were shown to be largely consistent with this argument.

From an applied point of view, our results suggest that mechanisms which implement a desirable outcome only under higher-order or even rational beliefs may fail in practice. In this sense, our findings are interesting for applied mechanism design since they indicate potential constraints on implementation due to bounded rationality.

Appendix

Derivation of RPI: The SPNE for the subgame of G_I in which players have already taken their insurance decision are immediate (cf. Table 1 in the text). Accordingly, we have the following pure strategy SPNE: no player being insured if $c \geq 11 + \varepsilon$ only player 1 being insured if $c \in (0, 9 + \varepsilon]$, only player 2 being insured if $c \in [\varepsilon, 11 + \varepsilon]$.

We now solve for a mixed strategy Nash equilibrium. From the above analysis, it follows that $c \in [\varepsilon, 9 + \varepsilon]$ is a necessary condition for a strictly mixed Nash equilibrium. Let p_i^* be the probability with which player i buys the insurance. In equilibrium, given p_j^* for $j \neq i$, player i is indifferent between buying the insurance and being uninsured. Hence, the expected payoff from not buying the insurance equals the expected payoff from buying the insurance. For player 1 the former is given by $17p_2^* + 8(1 - p_2^*)$, the latter by $17 + \varepsilon - c$. Using the above indifference condition and solving for p_2^*, we obtain

$$p_2^* = 1 + \frac{\varepsilon - c}{9}.$$

Similarly, we obtain

$$p_1^* = 1 - \frac{c}{11 + \varepsilon}.$$

Hence, there is a unique strictly mixed SPNE, in which player 1 buys the insurance with probability $p_1^* = 1 - \frac{c}{11 + \varepsilon}$ and player 2 buys the insurance with probability $p_2^* = 1 + \frac{\varepsilon - c}{9}$.

Accordingly, for $c \in (\varepsilon, 9 + \varepsilon)$, there are three SPNE, namely one mixed equilibrium in which both players are insured with the probabilities just derived, and two pure equilibria in which one player insures with probability 1 and the other does not insure. Equilibrium play of the centipede game is according to Table 1. $\square$

Derivation of RPB: Again the SPNE of the CP-game for given bonus offers are immediate (cf. Table 2 in the text). Consequently, we have the following pure strategy SPNE: no player offering the bonus if $b \geq 14$, only player 1 offering the bonus if $b \in (2, 12]$, only player 2 offering the bonus if $b \in [5, 14]$.

For $b \in (5, 12)$ there is an additional equilibrium in mixed strategies. Let p_i^* be the probability with which player i offers the bonus. In equilibrium, given p_j^* for $j \neq i$, player i is indifferent between offering the bonus and not offering the bonus. For player 1, offering the bonus gives an expected payoff of $20p_2^* + (20 - b)(1 - p_2^*)$, while not offering the bonus leads to an expected payoff of $(15 + b)p_2^* + 8(1 - p_2^*)$. Using the indifference condition, we obtain:

$$p_2^* = \frac{12 - b}{7}.$$

Similar considerations for player 2 yield

$$p_1^* = 1 - \frac{b}{14}.$$

$\square$

References

ANDREONI, J. (1995): "Warm-Glow Versus Cold-Prickle: The Effects of Positive and Negative Framing on Cooperation in Experiments," *Quarterly Journal of Economics*, 110, 1–21.

COSTA-GOMES, M., V. CRAWFORD, AND B. BROSETA (2001): "Cognition and Behavior in Normal-Form Games: An Experimental Study," *Econometrica*, 69, 1193–1235.

CRAWFORD, V., AND N. IRIBERRI (2007a): "Fatal Attraction: Salience, Naivity, and Sophistication in Experimental 'Hide and Seek' Games," *American Economic Review*, 97, forthcoming.

——— (2007b): "Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private Auctions?," *Econometrica*, 75, 1721–1770.

FISCHBACHER, U. (1999): "z-Tree. Toolbox for Readymade Economic Experiments," IEW Working Paper No. 21, University of Zurich.

GREINER, B. (2004): "The Online Recruitment System ORSEE 2.0 - A Guide for the Organization of Experiments in Economics," Working Paper Series in Economics 10, Department of Economics, University of Cologne.

HO, T., C. CAMERER, AND K. WEIGELT (1998): "Iterated Dominance and Iterated Best Response in Experimental "p-Beauty Contests"," *American Economic Review*, 88, 947–969.

KATOK, E., M. SEFTON, AND A. YAVAS (2002): "Implementation by Iterative Dominance and Backward Induction: An Experimental Comparison," *Journal of Economic Theory*, 104, 89–103.

KREPS, R., P. MILGROM, J. ROBERTS, AND R. WILSON (1982): "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory*, 27, 245–252.

MCKELVEY, R., AND T. PALFREY (1992): "An Experimental Study of the Centipede Game," *Econometrica*, 60, 803–836.

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

NAGEL, R., AND F. TANG (1998): "Experimental Results on the Centipede Game in Normal Form: An Investigation on Learning," *Journal of Mathematical Psychology*, 42, 356–384.

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

STAHL, D., AND P. WILSON (1994): "Experimental Evidence on Players' Models of Other Players," *Journal of Economic Behavior and Organization*, 25, 309–327.

——— (1995): "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 10, 218–254.