Economics Division
University of Southampton
Southampton SO17 1BJ, UK

Discussion Papers in
Economics and Econometrics

Title Aggregate Information and the Centipede Game: a Theoretical and Experimental Study.

By Zacharias Maniadis, University of Southampton

No. 1215

ISSN 0966-4246

Abstract

Theoretical models often neglect the importance of ex-post feedback for equilibrium behavior in games. We test for this importance in a set of new experiments of the centipede game with feedback about aggregate behavior in each round. To make formal predictions, we develop a simple model that applies the framework of Dekel, Fudenberg, and Levine (2004). Unlike many popular models, our model, which combines self-confirming equilibrium with non-selfish payoffs, predicts that aggregate feedback has equilibrium effects. Our subjects exhibit Nash behavior more often than subjects in previous experiments, and aggregate feedback causes even stronger convergence to Nash equilibrium. However, after slightly changing the payoff structure of the experimental game, the treatment effects of aggregate information often shift in the opposite direction. From the policy point of view, the experimental results show that whether aggregate information generates more trust and higher social payoffs depends on the details of the game.

Keywords: Centipede Game, Bayesian Games, Self-Confirming Equilibrium, Social Trust

JEL Classification Codes: C71, C73, C91

Aggregate Information and the Centipede Game: a Theoretical and Experimental Study

January 22, 2011

1 Introduction

Does aggregate feedback[1] affect long-run behavior in a dynamic game with large populations? The literature has shown that aggregate feedback may affect equilibrium play in "recurring games" with social learning (Jackson and Kalai, 1997) and also in games with psychological preferences (Battigalli and Dufwenberg, 2009) or preferences with conformity or reciprocity (Cialdini and Goldstein, 2004, Levine, 1998).[2] In this paper we wish to show, theoretically and experimentally, that aggregate feedback matters, even without social learning about the payoff function, and without conditional social preferences. We start from the observation that the heterogeneous self-confirming equilibrium beliefs introduced in Fudenberg and Levine (1993) have the property that they are not robust to aggregate information release. The reason is that some individuals may maintain wrong beliefs, which may be falsified when true aggregate behavior is revealed. Fudenberg and Levine (1997) argue that these heterogeneous beliefs might play an important role in subjects' behavior in experiments of extensive-form games. In this paper we formalize and experimentally test their informal prediction that aggregate information release will have strong effects on players' equilibrium behavior, focusing on the centipede game.

There are two reasons for the choice of the centipede game. First, the results from previous centipede game experiments have been among the most famous examples of failure of game-theoretic predictions. A series of experimental and theoretical papers has been inspired by the classical experimental study of McKelvey and Palfrey (1992) (henceforth referred to as MP). It is of interest to examine whether aggregate information brings subjects'

[1] We will use the terms "aggregate feedback" and "aggregate information" interchangeably, to denote information about behavior in a large number of different matches, with different individual players in each match.

[2] Conformity, reciprocity and psychological preferences are called "conditional social preferences". They make aggregate feedback relevant for equilibrium play, because they imply that individuals' behavior depends on their beliefs about the behavior of peers, or about other individuals' beliefs. In dynamic games with random matching, each individual repeatedly plays a fixed stage game against individuals that belong to the other population-roles. In the absence of aggregate information, long-run beliefs, regarding the behavior of peers or regarding others' beliefs, need not be correct. This is because there is no repeated interaction with peers (only with opponents) and no feedback about others' beliefs, so agents need not learn the relevant information, even in the long run. Aggregate feedback implicitly provides information about others' beliefs, and explicitly reveals peers' behavior. Thus, it can change individuals' equilibrium beliefs and their behavior.

behavior closer to Nash equilibrium with selfish payoffs, or not. Second, the centipede game that we use is a two-stage trust game, and it is worthwhile examining whether aggregate information can induce more trust and increase social payoffs in the long run, or not.[3] Since aggregate information is often in the hands of policy-makers, there are important lessons that might be learned for economic policy from this exercise. Previous experiments of the trust game and related games indicate that it is possible that aggregate information encourages trusting behavior and improves social payoffs.[4]

In the first part of the paper, we propose a model based on the framework of Dekel, Fudenberg, and Levine (2004) (henceforth referred to as DFL), introducing a small fraction of non-selfish individuals in the population of player 2's of the centipede game. The presence of non-selfish individuals, who pass in every opportunity, is necessary, since heterogeneous beliefs alone cannot explain why an individual would pass in the last decision node of the centipede game, "giving away money" to others.[5] The model has similarities with the model developed by MP to explain their data, but the solution concepts are motivated by learning theory, and differ from standard concepts for Bayesian games. Unlike sequential equilibrium, the concept of type-heterogeneous self-confirming equilibrium (THSCE) does not assume correct beliefs about opponents' play.

In the second part of the paper, we present several experiments of two-player, four-move centipede games, which are designed to test the effects of aggregate information release. In our experimental sessions, each subject is randomly assigned to one of two fixed groups, and interacts with each member of the other group exactly once. There are three forms of information feedback. The control treatments, where no aggregate feedback is provided, are called "no information" treatments. In the "full information" treatments, subjects can see the aggregate proportions of *Pass* and *Take*, in each of the four decision nodes of the game, in the immediately previous round. In the "partial information" treatments, each subject observes these proportions only for the decision nodes that belong to the other group.[6] Two payoff structures are used. In the "initial-payoff" treatments, we use the same payoff structure as in MP (Figure 1). In the "modified-payoff" treatments, we introduce a small but salient change in the payoff structure, such that the cost, for player 2, of choosing *Pass* in the last node is lower (Figure 2). The experimental and theoretical results and contributions are summarized as follows:

1. In the initial-payoff treatment with no aggregate feedback, subjects reach the Nash equilibrium outcome (of the game where all individuals are assumed to have selfish

[3]The importance of social capital in modern societies has been greatly underscored by sociologists and political scientists (Fukuyama, 2001; 2002).

[4]These experiments typically find positive effects of aggregate information on social payoffs. Berg, Dickhaut, and McCabe (1995) used an one-round trust game, and found that aggregate information about the past behavior of another group of subjects increased social payoffs. Ortmann, Fitzgerald, and Boeing (2000) showed that this result is robust to the way information is presented and to prompts for strategic reasoning. In the field experiment of Frey and Meier (2004), revealing information about the fraction of the total population that performed a certain charitable action tended to increase the frequency of this action.

[5]A non-trivial fraction of subjects in previous experiments, as well as in our experiment, exhibit this behavior.

[6]Each group corresponds to a player-role in the game. Each of the four decision nodes belongs to one of the two groups. This means that, with partial information, subjects in the role of player 2 observe aggregate information only in the nodes that belong to subjects in the role player 1, and vice versa.

payoffs) more often in our experiments than in MP. This might be due to the fact that our UCLA subjects, who are part of a relatively large subject pool, are less likely to behave cooperatively than MP's Caltech subjects, who belong to a small subject pool. These results show that there are limits to the robustness of the strong results of MP.

2. In the initial-payoff treatments with aggregate information, strong convergence to the Nash equilibrium outcome takes place, both under full and under partial information release. Aggregate information has a negative effect on trust and social payoffs. Moreover, subjects' behavior does not seem to depend much on whether own-group information is revealed or not, which is consistent with our model. These results indicate that frequent Nash play can be observed in centipede experiments, even without changing the structure of MP's game.[7]

3. The small alteration in payoffs causes a remarkably large change in the effects of aggregate information. In the modified-payoff treatments, convergence goes in the opposite direction (that of increasing payoffs), in two of the three experimental sessions with aggregate information. Thus, aggregate information seems to increase average payoffs in this case. Hence, the results show that the relationship between aggregate information and social payoffs is not fixed, even for similar games. The policy message is that policy-makers should reveal only "optimistic" information. Moreover, the results show that there does not seem to exist a stable relationship between aggregate feedback and the performance of Nash equilibrium and its refinements.[8]

4. The behavioral change, caused by the payoff modification, is partly due to the fact that aggregate information, after each round, introduces strong path dependence, which tends to magnify initial trends in behavior. Simple regression analysis shows that aggregate play depends much more on play in the previous period, when aggregate information is publicly revealed, than when it is not. In addition, repeated-game effects[9] may have also played a role in the behavioral change.

5. Unlike several standard models, our simple deterministic model has relatively good performance. In the environment with no aggregate feedback, there are multiple equilibria. The model specifies conditions that have to hold in any equilibrium distribution of moves in the four decision nodes, and these agree with the data. In the environment with aggregate feedback, the model has only two equilibrium outcomes. We propose

[7]As we shall see in the next section, several experiments of the centipede game found stronger support in favor of the Nash prediction than MP. However, they changed the basic structure of MP's extensive form, often confounding more than one such modification.

[8]Our paper thus contributes to the literature on the effects of aggregate information on the performance of theoretical predictions. Harrison and McCabe (1996), after observing convergence to the subgame-perfect equilibrium in their ultimatum game experiments, claimed that aggregate information serves as a surrogate to common belief in rationality. This general assertion does not find support in the rest of the literature. Dufwenberg and Gneezy (2002) examined experimental auctions, where information about the entire vector of bids led average bids away from Nash equilibrium. Furthermore, as we have already seen, in trust and public goods games, aggregate information tends to reduce the performance of equilibrium predictions (it increases cooperation).

[9]These effects occur when subjects realize that by passing now they might affect the future behavior of opponents, and thus they might not choose a myopic best response.

an equilibrium selection device, and experimental results are generally consistent with the predicted equilibrium outcome.

6. The predictive accuracy of our model is reasonable, given that it is a deterministic model of experimental subjects' behavior, and the fact that multiplicity of equilibria is a common feature of self-confirming equilibrium models. Our application is one of the first that combine self-confirming equilibrium with non-selfish preferences. It also examines the intuitively natural, but rarely studied observation environment, where actions, but not types, are observed in each round of play.

Section 2 introduces the centipede game, and explains how our paper relates to the relevant theoretical and experimental literature. In section 3, we present the general framework of our model, and we apply it to the centipede game, deriving exact predictions. Section 4 introduces our experiments. Section 5 presents the results of our five experimental treatments. Section 6 considers the performance of our model and alternative theories. Section 7 discusses the dynamics of play, and the possibility of repeated-game effects. In section 8 we conclude, and point out possible directions for future research.

2 The centipede game and related literature

The behavior of experimental subjects in the centipede game has been one of the most intriguing results of the experimental economics literature. In the two-player centipede game (see Figure 1 for an example), two players share a monetary amount split into a large and a small pile, in a predetermined way for each terminal node. In each decision node, the player who moves can either "take" the large pile of money and the game ends, or "pass" and let the total amount multiply in size. A player should choose *Take* now, if he expects that the other player will choose *Take* in the subsequent move, but the player is better off choosing *Pass* now, if he expects that the other player will also choose *Pass* in the subsequent move. In its finite version, the centipede game has an obvious candidate for a prediction of how it will be played: the unique Nash equilibrium outcome (and of course the unique subgame-perfect equilibrium outcome), has player 1 choose *Take*, in the first move, with probability one.[10]

Early experimental studies found little support in favor of the Nash prediction. We shall call "equilibrium percentage" the percentage of experimental matches, in late rounds of play, which finish in the first terminal node (which corresponds to the Nash outcome). In MP, subjects' equilibrium percentage was no more than 8%. Nagel and Tang (1998), using the normal form of a 12-move centipede game, found an equilibrium percentage not exceeding 0.5%. Fey, McKelvey, and Palfrey (1996) used a game with constant social payoffs,[11] and the equilibrium percentage in their experiments was between 20% and 70%. In Rapoport et al.

[10] Any strategy where player 1 passes in his fist decision node with positive probability can be part of an equilibrium only if player 2 also passes in his first decision node with positive probability. But for this to be optimal, player 1 should pass with positive probability in his second decision node, and so on. Finally, player 2 should pass with positive probability in his last decision node, which is clearly not optimal given that this node is reached with positive probability.

[11] In this game, the predictions of Nash, fairness and focal point theories agree in the same predicted outcome (the first terminal node with probability one).

Figure 1: A Two–Player, Four-Move Centipede Game with Geometrically Increasing Payoffs

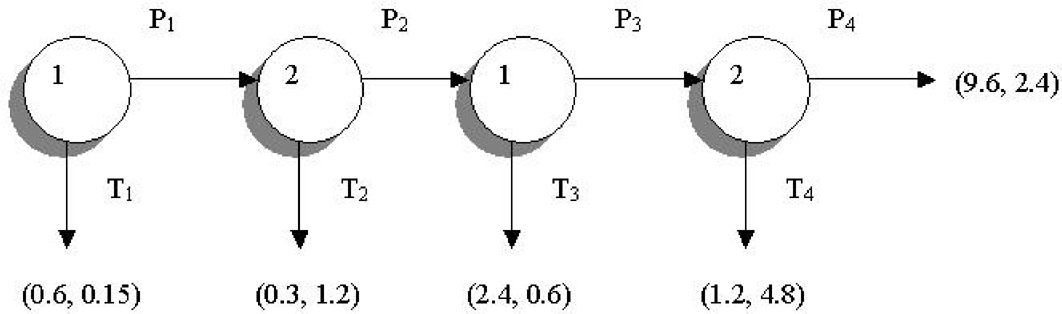

(2003), the equilibrium percentage was between 30% and 40%, in an experiment where each "inning" of choices involved three players, rather than two.[12] Murphy, Rapoport, and Parco (2006) used a continuous-time, symmetric version of the centipede game, and observed high equilibrium rates, especially for large groups of subjects. Finally, in Palacios-Huerta and Volij (2009), subjects included sophisticated chess players, who tended to play according to the Nash equilibrium prediction with higher frequency.[13]

Many different theoretical explanations have been offered for these results. MP introduced an incomplete information game with a small fraction of altruists, who pass in every node. McKelvey and Palfrey (1998) later developed the concept of Agent Quantal Response Equilibrium (AQRE), which captures subjects' behavior in their original centipede experiment, as well as results from other games. This concept assumes that agents have correct beliefs about others' play, but imperfectly best-respond to these beliefs. Zauner (1999) used a similar idea, but he employed a different structure of the noise, and also interpreted the noise differently. We have already seen that Fudenberg and Levine (1997) explained subjects' behavior as best-responding to beliefs, which need not be correct, but only consistent with the evidence. We enrich Fudenberg and Levine's explanation by adding some individuals with non-standard preferences, which is necessary in order to fully explain the results.[14]

3 Our model

We shall first motivate and introduce the general framework, which we will use for modeling the long-run behavior of our subjects. Subsequently, we shall apply this framework to the particular extensive forms that we use in the experiments, and we will get exact predictions.

As we shall explain in section 6, models that use standard solution concepts, such as Nash equilibrium, AQRE, and sequential equilibrium, are not appropriate for capturing the

[12] In addition to this modification, stakes were much higher, the number of rounds was 60, rather than 10, and the last terminal node gave zero payoffs to all players.

[13] For comparison, in our control treatment, which was essentially a replication of MP, the equilibrium percentage was about 30%. In our "full information" treatment with the initial payoffs, with no changes in the basic game structure of MP, the equilibrium percentage was 50%.

[14] The level-k analysis of Kawagoe and Takizawa (2008) also explains the results of experiments of the centipede game using the possibility of wrong beliefs. However, they focus on initial rounds only, whereas we are interested in equilibrium play.

Figure 2: Our Two–Player Centipede Game with Modified Payoffs

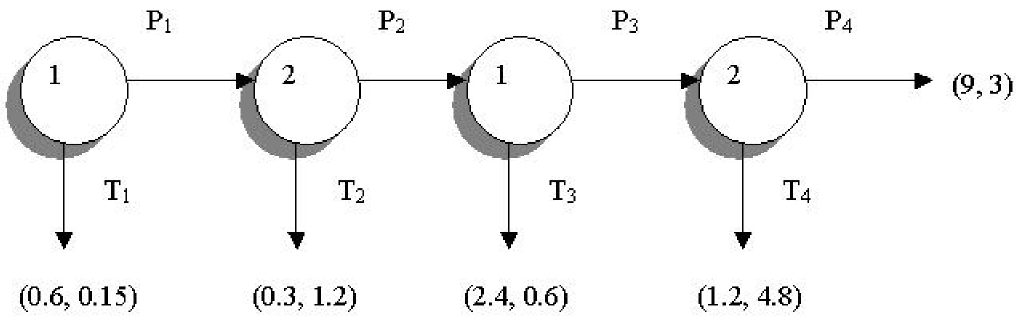

equilibrium effects of aggregate information. The reason is that they tend to ignore factors such as the ex-post feedback from play (Armantier, 2004, p. 238). In our experiments we are interested in the long-run behavior of subjects, in a setting of repeated interactions with anonymous matching. The appropriate steady-state concepts for capturing this behavior are developed in learning-motivated equilibrium models.[15] A well-known result from these models is that the steady state of a belief-based learning process is a "conjectural" or "self-confirming" equilibrium (Battigalli, 1987, Fudenberg and Levine, 1993, Kalai and Lehrer, 1993). In such an equilibrium, each individual best-responds to beliefs (about opponents' strategies) that need not be correct, but only consistent with the feedback that the agent receives, given her steady-state behavior.

In the self-confirming equilibrium (SCE) tradition, the information feedback, which players receive each time the game is played, is important for determining long-run behavior. Any information about opponents' play, that an agent eventually learns from personal experience, is contained in the aggregate feedback. But the converse is not true. Therefore, aggregate information about play has a natural effect on equilibrium beliefs and behavior. Each individual's beliefs should respect this information, so the set of possible beliefs in a self-confirming equilibrium is typically smaller when aggregate information is revealed.

The general framework that we shall use in our model follows DFL, who consider the implications of learning theory for Bayesian games.[16] We shall model the interaction between subjects as a Bayesian game (Figure 3), because we introduce a non-selfish type. We chose this modeling approach, because we believe that the introduction of non-selfish individuals, whose behavior is very basic, generates the simplest self-confirming equilibrium model, capable of fully capturing the data of MP. In addition, using the DFL framework, we can

[15]We should note here that an important underlying assumption in the "learning in games" literature is that individuals' beliefs are formed purely by experience. No prior knowledge of opponents' payoffs is assumed, and, in general, agents do not deduct behavior from prior information about the parameters of the game.

[16]Our experimental interactions can only be imperfectly captured by the environment of Fudenberg and Levine (1993), where individuals interact anonymously in an extensive-form game. We adhere to some of the assumptions of this model: each individual knows the extensive form of the game and her payoffs for each terminal node, but not the payoffs of other individuals. However, we need to depart from the assumption that individuals observe the realized terminal node of their game after each match. The reason is that we believe that it does not correspond with the experimental environment. In the experiments, subjects only observe the moves of opponents on the path of play, but not their types.

Figure 3: The Incomplete Information Centipede Game

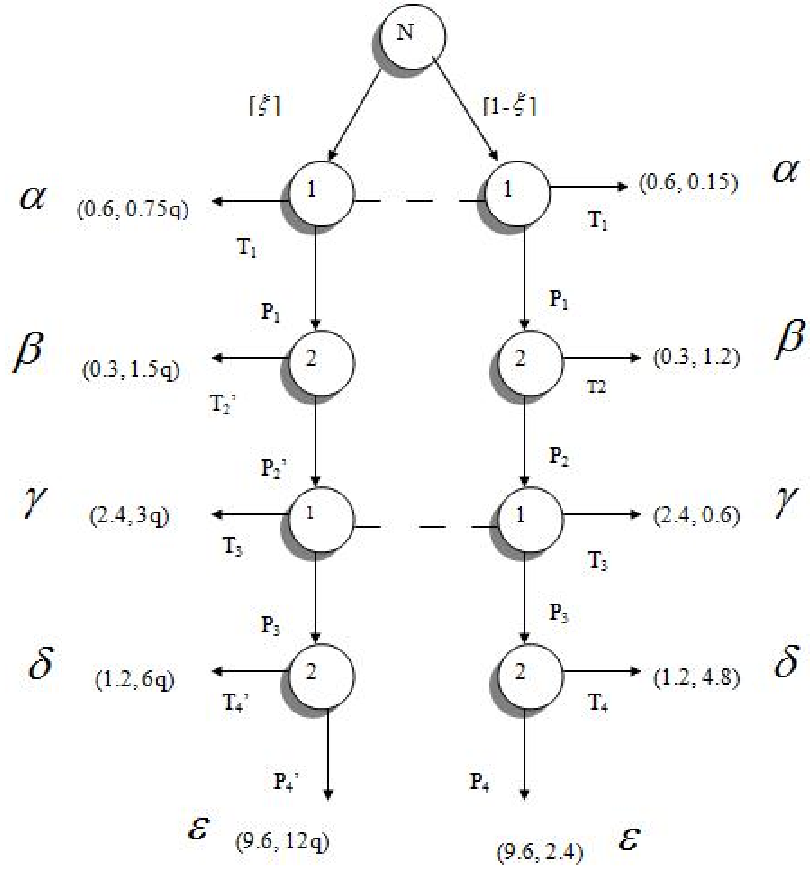

model aggregate information release as a change in the "signal function", which determines the feedback, which each individual receives in each period.

We will also make the assumption that an individual always has the same type (rather than the type being drawn independently in each period).[17] This is in accordance with the observation made in MP, that in their data, there is a nontrivial fraction of subjects, who always choose *Pass*. Therefore, the framework that we shall now consider corresponds to "fixed types for each agent, but diversity across agents in the same role" (DFL, p. 298).

3.1 The formal framework

The framework considers static, simultaneous-move games with I player-roles. All of the parameters of the game are assumed to be finite. The characteristics of the static stage game are as follows. Player i's type is denoted $\theta_i \in \Theta_i$. Players simultaneously choose actions, which we denote $a_i \in A_i$, for each player i. A strategy for player i, σ_i, is a function from own types to probability distributions over actions, so that $\sigma_i(\theta_i) \in \Delta(A_i)$.[18] The set

[17]Although an agent will have a single type in all periods, because of random matching, her opponents in each match shall not know the agent's type.

[18]We need to emphasize our specific interpretation of σ_i. We assume that each individual in population i has a fixed type, and chooses a pure action in A_i. However, the population as a whole randomizes across

of all such strategies for player i is denoted Σ_i, and $\sigma_i(a_i|\theta_i)$ denotes the probability that $\sigma_i(\theta_i)$ assigns to action a_i. Action, type and strategy profiles are denoted $a \in A$, $\theta \in \Theta$, and $\sigma \in \Sigma$, respectively. Player i's payoffs $u_i(a, \theta)$ depend on the profile of actions played and the profile of types. If the game has private values, each player's utility depends only on the profile of actions, and her own type: $u_i(a, \theta) = u_i(a, \theta_i)$.

DFL's solution concept is motivated by the idea that the above stage game is played repeatedly, but anonymously, by individuals that belong to large populations. Individuals' types are drawn, before the first interaction, according to a probability distribution p. After the type of each individual is drawn, individuals are randomly matched in each period to play the fixed stage game, with a matching process independent of the players' types. The matching process ensures that the types of individuals in a given match are independent. Individuals do not know opponents' strategies and the actual distribution of types, by player i has "conjectures" $\widehat{\sigma}_{-i} \in \prod_{j \neq i} \Sigma_j$ about opponents' strategies, and "interim beliefs" $\widetilde{\mu}^{\theta_i} \in \Delta(\Theta_{-i})$ about opponents' types. Players update their beliefs and conjectures, after getting feedback on opponents' moves and types, each time they play the game. The feedback received after each interaction is described by the private deterministic signal function (also called "observation function") $y_i(a, \theta)$.[19]

DFL's focus is on the possible steady states of the dynamic process that may arise, starting from any priors about opponents' types and strategies. The following solution concept captures the fact that different individuals, in a given population, may have different personal experience stemming from equilibrium behavior, if their types or their equilibrium actions differ. Accordingly, their beliefs and conjectures, which must be consistent with this personal experience, may differ in equilibrium. Equilibrium actions are only required to be optimal given beliefs and conjectures.

Definition 1 (DFL p. 298). A strategy profile σ is a type-heterogeneous self-confirming equilibrium if, for each player i, and for each $\widehat{a}_i$ and θ_i such that $p(\theta_i)\sigma_i(\widehat{a}_i|\theta_i) > 0$, there are conjectures $\widehat{\sigma}_{-i}$ and interim beliefs $\widetilde{\mu}^{\theta_i}$ (both of which can depend on $\widehat{a}_i$ and θ_i), such that:

1. $\widehat{a}_i \in argmax_{a_i} \sum_{a_{-i}, \theta_{-i}} u_i(a_i, a_{-i}, \theta_i, \theta_{-i}) \widetilde{\mu}^{\theta_i}(\theta_{-i}) \widehat{\sigma}_{-i}(a_{-i}|\theta_{-i})$

2. For all $\widetilde{y}_i$ in the range of y_i, the following holds:

$$\sum_{\{a_{-i}, \theta_{-i} : y_i(\widehat{a}_i, a_{-i}, \theta_i, \theta_{-i}) = \widetilde{y}_i\}} \widetilde{\mu}^{\theta_i}(\theta_{-i}) \widehat{\sigma}_{-i}(a_{-i}|\theta_{-i}) = \sum_{\{a_{-i}, \theta_{-i} : y_i(\widehat{a}_i, a_{-i}, \theta_i, \theta_{-i}) = \widetilde{y}_i\}} p(\theta_{-i}) \sigma_{-i}(a_{-i}|\theta_{-i})$$

Condition 2 simply says that the probability distribution of signals, induced by player i's beliefs and conjectures, coincides with the distribution of signals induced by the actual distribution of types and the actual strategies of opponents.

types and actions. Each population i is partitioned into subpopulations, corresponding to different types. A strategy σ_i, as defined above, further partitions each such subpopulation into sets of individuals that choose the same action.

[19]Note that this function is the same for all individuals in a given population, but the actual feedback that each individual receives depends on her action and type. Therefore, the feedback from equilibrium behavior that individuals receive can differ.

Distributions of actions of each population

As we shall see, since our game has private values, players simply best-respond to their beliefs about the distribution of actions of other populations. It will therefore be convenient to introduce some new notation, in order to describe these distributions.

Let the marginal of p on i's types be denoted $p_i \in \Delta(\Theta_i)$. Population i's distribution of types p_i and its strategy σ_i determine the distribution of actions of this population as follows. Consider a specific action a_i. The fraction of population i that chooses a_i is simply given by the sum (over types θ_i) of the probabilities assigned to action a_i by $\sigma_i(\theta_i)$, each probability weighted by $p_i(\theta_i)$, the fraction of population i that has type θ_i. To formalize this, we define, for each i, the function $s_i : \Sigma_i \times \Delta(\Theta_i) \to \Delta(A_i)$ as follows. For each $a_i \in A_i$, $s_i(a_i|\sigma_i, p_i) \equiv \sum_{\theta_i \in \Theta_i} p_i(\theta_i)\sigma_i(a_i|\theta_i)$, where $s_i(a_i|\sigma_i, p_i)$ denotes the probability that $s_i(\sigma_i, p_i)$ assigns to action a_i. So, $s_i(\sigma_i, p_i)$ is simply the distribution of actions of population i induced by σ_i and p_i.

Let an individual player in population j have interim beliefs $\widetilde{\mu}^{\theta_j}$ and conjectures $\widehat{\sigma}_{-j} = (\widehat{\sigma}_1, \widehat{\sigma}_2, \ldots, \widehat{\sigma}_{j-1}, \widehat{\sigma}_{j+1}, \ldots, \widehat{\sigma}_I)$. Denote the marginal of the beliefs $\widetilde{\mu}^{\theta_j}$ on i's types by $\widetilde{\mu}_i^{\theta_j}$. Then, the distribution $s_i(\widehat{\sigma}_i, \widetilde{\mu}_i^{\theta_j})$ can be interpreted as the distribution of actions of population i, which this individual player j expects, given her interim beliefs $\widetilde{\mu}^{\theta_j}$ and her conjectures $\widehat{\sigma}_{-j}$. We can thus define the profile of expected distributions of actions for players other than j, induced by j's beliefs $\widetilde{\mu}^{\theta_j}$ and conjectures $\widehat{\sigma}_{-j}$, as follows: $s_{-j}(\widehat{\sigma}_{-j}, \widetilde{\mu}^{\theta_j}) = [s_1(\widehat{\sigma}_1, \widetilde{\mu}_1^{\theta_j}), \ldots, s_{j-1}(\widehat{\sigma}_{j-1}, \widetilde{\mu}_{j-1}^{\theta_j}), s_{j+1}(\widehat{\sigma}_{j+1}, \widetilde{\mu}_{j+1}^{\theta_j}), \ldots, s_I(\widehat{\sigma}_I, \widetilde{\mu}_I^{\theta_j})]$.

Aggregate information about distributions of actions

The private signal functions of DFL correspond to the feedback that individuals receive by personal experience alone. So, they are not appropriate for describing aggregate information revelation. It is important to emphasize that when aggregate information is released, the signal that an individual player i observes depends on the strategy profile σ and the distribution of types p. That is, the signal depends on the aggregate behavior and the aggregate distribution of types, rather than the actions and types of players in individual i's own matches. Hence, to capture aggregate information release, we need a signal function Y_i that has as its arguments σ and p. With such a signal function, two different individuals who belong to population i receive exactly the same signal each period, regardless of their own behavior, and the behavior of the individuals with whom they are matched.

In our experiments, individuals receive aggregate feedback about the distributions of actions of opponent populations (not about opponents' strategies or types). In particular, subjects from population i are informed that the profile of distributions of actions, of other populations, lies in a certain subset of the set $\prod\limits_{j \neq i} \Delta(A_j)$, the set of possible distribution profiles. So, the observation function of population i maps the profile of strategies and the joint distribution of types into a set of profiles of opponents' action distributions. Hence, $Y_i : \Sigma \times \Delta(\Theta) \to \mathcal{P}(\prod\limits_{j \neq i} \Delta(A_j))$, where the notation $\mathcal{P}(X)$ denotes the power set of X. In other words, $Y_i(\sigma, p)$ is the set that contains all the profiles of distributions of actions for players other than i, consistent with the aggregate feedback that player i receives when the profile of strategies is σ and the probability distribution of types is p. Notice that the exact form of Y_i depends on the structure of the game, and we shall soon see this form for our centipede game. The equilibrium notion needs to be slightly modified, when aggregate information about the distributions of actions is provided.

Definition 2. A strategy profile σ is a type-heterogeneous self-confirming equilibrium with aggregate information if, for each player i, and for each $\widehat{a_i}$ and θ_i such that $p(\theta_i)\sigma_i(\widehat{a}_i|\theta_i) > 0$, there are conjectures $\widehat{\sigma}_{-i}$ and interim beliefs $\widetilde{\mu}^{\theta_i}$(both of which can depend on $\widehat{a_i}$ and θ_i), such that:

1. $\widehat{a}_i \in argmax_{a_i} \sum_{a_{-i},\theta_{-i}} u_i(a_i, a_{-i}, \theta_i, \theta_{-i})\widetilde{\mu}^{\theta_i}(\theta_{-i})\widehat{\sigma}_{-i}(a_{-i}|\theta_{-i})$

2. $s_{-i}(\widehat{\sigma}_{-i}, \widetilde{\mu}^{\theta_i}) \in Y_i(\sigma, p)$

Condition 2 says that the expected profile of distributions of actions of opponents, induced by conjectures and interim beliefs, must be consistent with the aggregate feedback induced by the true strategies and distribution of types.

3.2 The centipede game in the DFL framework

Consider the one-sided incomplete information model of the centipede game of Figure 3. Following MP, we assume that there is a small fraction of individuals with non-selfish preferences. However, unlike MP, non-selfish agents exist only in population 2. Non-selfish (we shall also call them "altruistic") individuals get payoffs commensurate to the total monetary payoffs of a terminal node ($q > 0$). The fraction of individuals in population 2, who have altruistic preferences, is given by the parameter ξ, where we assume that $1/30 < \xi < 1/10$. We shall analyze the steady states of a system with large populations, who repeatedly play this game, with random matching. Before the first interaction, individuals in population 2 learn their type.

As we shall now show, the DFL framework can easily accommodate this extensive-form game by defining the appropriate action sets and signal functions. With respect to the action sets, first consider the reduced normal form of the complete-information centipede game of Figure 1. Each action in the model of DFL will correspond to a pure reduced normal form strategy. Accordingly, the set of players is $I = \{1, 2\}$ and the action sets for players 1 and 2 are $A_1 = \{T_1, P_1T_3, P_1P_3\}$ and $A_2 = \{T_2, P_2T_4, P_2P_4\}$. For example, the action P_1T_3 for player 1 corresponds to the reduced normal form strategy "choose *Pass* in decision node 1 and choose *Take* in decision node 3" (the second decision node of player 1). Player 1 has one possible type, and player 2 has two possible types, so we set $\Theta_1 = \{\theta_1\}$ and $\Theta_2 = \{\theta_2^1, \theta_2^2\}$, where θ_2^1 is the selfish type. Since there is only one possible type for player 1, the interim beliefs of player 2 play no role, and they will be omitted in the following analysis. Moreover, the interim beliefs of the unique type of player 1 will be denoted $\widetilde{\mu}_1$ for notational simplicity. The probability distribution over types, p, simply assigns probability ξ to $\{\theta_1, \theta_2^2\}$ and probability $(1 - \xi)$ to $\{\theta_1, \theta_2^1\}$. Clearly, the game has private values.

A strategy for player 1 is a distribution over actions, so $\sigma_1 \in \Delta(A_1)$. For player 2, a strategy is a function from $\{\theta_2^1, \theta_2^2\}$ to the set $\Delta(A_2)$. The functions $s_1 : \Sigma_1 \to \Delta(A_1)$

and $s_2 : \Sigma_2 \times \Delta(\Theta_2) \to \Delta(A_2)$ describe population distributions of actions induced by strategies and distributions of types, as explained in section 3.1. For example, $s_2(T_2|\sigma_2, p) = (1-\xi) \cdot \sigma_2(T_2|\theta_2^1) + \xi \cdot \sigma_2(T_2|\theta_2^2)$ is the fraction of population 2 that chooses action T_2, if the strategy of this population is σ_2 and the distribution of types is given by p. Notice that $s_1(\sigma_1) = \sigma_1$ for all $\sigma_1 \in \Sigma_1$.

The most important part in analyzing the incomplete information centipede game within the DFL framework is defining the appropriate signal functions. These functions will be constructed based on the implications of the extensive form of Figure 3, for the feedback that individuals will receive each time they play the game. Importantly, the signal functions capture all the relevant aspects of the extensive form for our purposes.

Notice that the signal functions will depend on whether or not aggregate feedback is provided. Accordingly, we shall first define them for the setting with no aggregate feedback, and we will also state some conditions that need to hold in a THSCE without aggregate information. Subsequently, we shall specify the appropriate signal functions in the presence of aggregate feedback about the distributions of actions of others.

3.2.1 The signal functions in the setting without aggregate information

First consider the setting without aggregate information. As we have noted, each time the centipede game is played, subjects observe opponents' moves along the path of play, but not opponents' types. Hence, the most natural assumption is that players do not observe the terminal node of the "grand" incomplete information game. A useful way to think about what they actually observe is to consider any subgame starting from any initial move by Nature. A player observes which terminal node of this subgame occurs, without being able to distinguish between subgames. In Figure 3, we label the terminal nodes of each such subgame, using the letters $\{\alpha, \beta, \gamma, \delta, \epsilon\}$. The observation function y (same for the two players) can be defined as follows:

1. $y(T_1, T_2, \theta_2^1) = y(T_1, T_2, \theta_2^2) = y(T_1, P_2T_4, \theta_2^1) = y(T_1, P_2T_4, \theta_2^2) = y(T_1, P_2P_4, \theta_2^1) = y(T_1, P_2P_4, \theta_2^2) = \alpha$
2. $y(P_1T_3, T_2, \theta_2^1) = y(P_1T_3, T_2, \theta_2^2) = y(P_1P_3, T_2, \theta_2^1) = y(P_1P_3, T_2, \theta_2^2) = \beta$
3. $y(P_1T_3, P_2P_4, \theta_2^1) = y(P_1T_3, P_2P_4, \theta_2^2) = y(P_1T_3, P_2T_4, \theta_2^1) = y(P_1T_3, P_2T_4, \theta_2^2) = \gamma$
4. $y(P_1P_3, P_2T_4, \theta_2^1) = y(P_1P_3, P_2T_4, \theta_2^2) = \delta$
5. $y(P_1P_3, P_2P_4, \theta_2^1) = y(P_1P_3, P_2P_4, \theta_2^2) = \epsilon$

The information (regarding opponent's chosen action) that an agent can infer from an interaction, depends on both her behavior and her opponent's behavior. For example, if, in a given interaction, player 1 chooses T_1, then she observes nothing about player 2's action. Hence, all actions of player 2 should be viewed as possible. The first condition in the signal function says exactly this: when action T_1 is played, the observation function does not distinguish between 2's actions. On the other hand, when player 1 chooses P_1P_3, all information sets of player 2 can be reached, and a different signal is observed for each action of player 2 (see conditions $2, 4$ and 5). Hence, player 1 can fully distinguish between the actions of player 2. Finally, when player 1 chooses P_1T_3, she can tell whether player 2 has chosen T_2, but she cannot distinguish between P_2T_4 and P_2P_4, as indicated in condition 3.

From the perspective of player 2, when she chooses action T_2 she may only observe the behavior of player 1 in the latter's first information set. Accordingly, player 2 may distinguish T_1 from the other two actions of player 1, since when she observes "α" she infers T_1. However, she cannot distinguish between P_1T_3 and P_1P_3, since when she observes "β" she does not know which of the two actions player 1 has chosen. When player 2 chooses any of the other two actions, P_2T_4 or P_2P_4, she can distinguish between any of player 1's actions. This is because a unique action of player 1 is compatible with each signal, given player 2's own action. Note that neither player 1 nor player 2 can distinguish between the opponent's types, for any profile of actions and types.

Type-heterogeneous self-confirming equilibrium with no aggregate information

We shall use a simple lemma to illustrate some conditions that need to hold in a THSCE of our incomplete information four-move centipede game. To make the conditions easier to follow, we shall use two general payoff functions $u_1(a_1, a_2)$, $u_2(a_2, a_1, \theta_2^1)$ (that satisfy the condition of private values), and we shall also use the general distribution of types p. Private values simplify the THSCE conditions significantly, since each player's strategy is a best response to her expected distribution of opponent's actions. In particular, the expected utility of an action, given the strategy and the distribution of types of the other population, can be written in the following simplified way, for player 1 and selfish player 2:

$$Eu_1(a_1, \sigma_2, p) = u_1(a_1, T_2) \cdot s_2(T_2|\sigma_2, p) + u_1(a_1, P_2T_4) \cdot s_2(P_2T_4|\sigma_2, p) + u_1(a_1, P_2P_4) \cdot s_2(P_2P_4|\sigma_2, p)$$

$$Eu_2(a_2, \sigma_1) = u_2(a_2, T_1, \theta_2^1) \cdot s_1(T_1|\sigma_1) + u_2(a_2, P_1T_3, \theta_2^1) \cdot s_1(P_1T_3|\sigma_1) + u_2(a_2, P_1P_3, \theta_2^1) \cdot s_1(P_1P_3|\sigma_1)$$

Lemma 1: The definition of type-heterogeneous self-confirming equilibrium implies the following conditions for our game:

1. If σ is a THSCE such that $\sigma_1(T_1) > 0$, there are conjectures and beliefs $\widehat{\sigma}_2, \widetilde{\mu}_1$ such that T_1 maximizes $Eu_1(a_1, \widehat{\sigma}_2, \widetilde{\mu}_1)$ with respect to a_1.[20]

2. If σ is a THSCE such that $\sigma_1(P_1T_3) > 0$, there are conjectures and beliefs $\widehat{\sigma}_2, \widetilde{\mu}_1$ such that P_1T_3 maximizes $Eu_1(a_1, \widehat{\sigma}_2, \widetilde{\mu}_1)$ with respect to a_1, and also such that $s_2(T_2|\widehat{\sigma}_2, \widetilde{\mu}_1) = s_2(T_2|\sigma_2, p)$.

3. If σ is a THSCE such that $\sigma_1(P_1P_3) > 0$, there are conjectures and beliefs $\widehat{\sigma}_2, \widetilde{\mu}_1$ such that P_1P_3 maximizes $Eu_1(a_1, \widehat{\sigma}_2, \widetilde{\mu}_1)$ with respect to a_1, and also such that $s_2(T_2|\widehat{\sigma}_2, \widetilde{\mu}_1) = s_2(T_2|\sigma_2, p)$, $s_2(P_2T_4|\widehat{\sigma}_2, \widetilde{\mu}_1) = s_2(P_2T_4|\sigma_2, p)$ and $s_2(P_2P_4|\widehat{\sigma}_2, \widetilde{\mu}_1) = s_2(P_2P_4|\sigma_2, p)$.

[20] Notice that condition (2) in Definition 1 has no bite in this case, so there are no restrictions on possible beliefs and conjectures.

4. If σ is a THSCE such that $\sigma_2(T_2|\theta_2^1) > 0$, there are conjectures $\widehat{\sigma}_1$ such that T_2 maximizes $Eu_2(a_2, \widehat{\sigma}_1)$ with respect to a_2, and also such that $s_1(T_1|\widehat{\sigma}_1) = s_1(T_1|\sigma_1)$.

5. If σ is a THSCE such that $\sigma_2(P_2T_4|\theta_2^1) > 0$, there are conjectures $\widehat{\sigma}_1$ such that P_2T_4 maximizes $Eu_2(a_2, \widehat{\sigma}_1)$ with respect to a_2, and also such that $s_1(T_1|\widehat{\sigma}_1) = s_1(T_1|\sigma_1)$, $s_1(P_1T_3|\widehat{\sigma}_1) = s_1(P_1T_3|\sigma_1)$ and $s_1(P_1P_3|\widehat{\sigma}_1) = s_1(P_1P_3|\sigma_1)$.

6. If σ is a THSCE such that $\sigma_2(P_2P_4|\theta_2^1) > 0$, there are conjectures $\widehat{\sigma}_1$ such that P_2P_4 maximizes $Eu_2(a_2, \widehat{\sigma}_1)$ with respect to a_2, and also such that $s_1(T_1|\widehat{\sigma}_1) = s_1(T_1|\sigma_1)$, $s_1(P_1T_3|\widehat{\sigma}_1) = s_1(P_1T_3|\sigma_1)$ and $s_1(P_1P_3|\widehat{\sigma}_1) = s_1(P_1P_3|\sigma_1)$.

The proof simply applies the definition of THSCE, and uses the compact notation for the distributions of actions in order to simplify the expressions. Notice that the outcome where $\sigma_1(T_1) = 1$ is clearly the outcome of a THSCE, since action T_1 can be rationalized by unconstrained beliefs (condition 1). This is true both in the presence and the absence of aggregate information. The consistency conditions on beliefs can be easily understood if one considers the fact that each individual uses a pure action in equilibrium. If the action that the individual chooses is sufficiently informative, the distribution of opponents' actions that she expects equals the true distribution. The informativeness of an action depends on the information sets of opponents (in the original extensive form of the game) that it reaches with positive probability.

3.2.2 The signal functions in the setting with aggregate information

Recall that with aggregate information, all individuals in one population-role observe a set of possible distributions of actions for the other population-role. This is the aggregate signal. We need to emphasize the importance of the fact that the centipede game is an extensive-form game. Because of this fact, the aggregate signal may not be fully informative (meaning that it may not be a point in the space of the distributions of actions of the other player). The aggregate signal received by population i depends on which information sets of the other population (again, in the original extensive form of the game) are reachable[21] given σ_i and p_i.

Remember that a subject, in our experiments with aggregate feedback, observes the aggregate proportions of *Pass* and *Take* in each of the two decision nodes of the opponent, but only if such a proportion is available. Whether this proportion is available or not, depends on the behavior of the subject's own population. Assume that the behavior of population 1 is such that action P_1P_3 is never played, but P_1T_3 is played with positive probability. Then, the first information set of player 2 is reachable (and is actually reached with positive probability), but the second information set of player 2 is not reachable.[22] Thus, the information that individuals in population 1 receive, by observing the aggregate signal, is the fraction of individual 2's that play T_2 (remember that this fraction is determined

[21] An information set is reachable, given population i's behavior and distribution of types, if it can be reached with positive probability for some strategy profile of i's opponents. Whether this set is actually reached with positive probability, or not, also depends on the behavior of i's opponents.

[22] In other words, the proportions of *Pass* and *Take*, in the second decision node of population 2, are not available for subjects to see.

by 2's strategy σ_2 and the distribution of types p). On the other hand, if population 1's strategy is such that P_1P_3 is played with positive probability, both information sets of player 2 are reachable. Hence, the aggregate signal will be a point in the space of distributions of player 2's actions.[23] Finally, if all individual 1's play T_1, then no information set of player 2 is reachable, and the aggregate signal provides no information (it is the whole space). Accordingly, the observation function for population 1 is:

- $Y_1(\sigma_1, \sigma_2, p) = \Delta(A_2)$, if $s_1(T_1|\sigma_1) = 1$
- $Y_1(\sigma_1, \sigma_2, p) = \{f \in \Delta(A_2) : f(T_2) = s_2(T_2|\sigma_2, p)\}$, if $s_1(T_1|\sigma_1) \neq 1, s_1(P_1P_3|\sigma_1) = 0$
- $Y_1(\sigma_1, \sigma_2, p) = \{f \in \Delta(A_2) : f(T_2) = s_2(T_2|\sigma_2, p), f(P_2T_4) = s_2(P_2T_4|\sigma_2, p), f(P_2P_4) = s_2(P_2P_4|\sigma_2, p)\}$, if $s_1(P_1P_3|\sigma_1) \neq 0$

With respect to the aggregate signal function of population 2, notice that the first information set of player 1 is always reached, regardless of 2's actions. If the behavior of population 2 is such that all individuals choose action T_2, then the aggregate signal reveals only the fraction of population 1 that chooses T_1. If there exists a positive fraction of population 2 that plays P_2T_4 or P_2P_4, then both information sets of player 1 are reachable. Then, the aggregate signal reveals the fraction of population 1 that plays each of the three actions.[24] So, for population 2, the observation function is the following:

- $Y_2(\sigma_1, \sigma_2, p) = \{g \in \Delta(A_1) : g(T_1) = s_1(T_1|\sigma_1)\}$, if $s_2(T_2|\sigma_2, p) = 1$
- $Y_2(\sigma_1, \sigma_2, p) = \{g \in \Delta(A_1) : g(T_1) = s_1(T_1|\sigma_1), g(P_1T_3) = s_1(P_1T_3|\sigma_1), g(P_1P_3) = s_1(P_1P_3|\sigma_1)\}$, if $s_2(T_2|\sigma_2, p) \neq 1$

Consider equilibria with aggregate information, as specified in Definition 2. Condition 2 of the definition implies that each individual's beliefs and conjectures, in population i, is such that the distribution of opponents' actions, induced by these beliefs and conjectures, lies in $Y_i(\sigma_i, \sigma_j, p)$. This captures the fact that each individual's beliefs about the opponents' distribution of actions respects aggregate information.

3.3 Characterizing possible equilibrium outcomes

We are interested in making predictions about subjects' behavior in experiments, and in comparing the predictive performance of our model with the performance of other models.

[23] The second information set of player 2 will not be reached if T_2 is chosen by all individual 2's. However, in this case, the information from the behavior in the first information set of player 2 is enough to pinpoint a unique distribution of actions for population 2.

[24] It is important to underline a significant assumption that we will be using in the analysis. Because of the fact that there are two large populations of individual players, we consider a small fraction of one population as representative of the population. For example, even if only 2% of player 1's choose *Pass* in information set 1, we consider the behavior of the individual 2's, with whom they are matched, as representative of population 2. Since the interaction is assumed to be completely anonymous, these individual 2's have been randomly chosen from a very large population. Hence, these individuals are a large enough random sample, such that their behavior mirrors the strategy of the whole population 2.

In the experiments we do not observe subjects' reduced normal form strategies (actions), but only their choices on the path of play. Moreover, the predictions of the AQRE model are presented as distributions of moves in each decision node of the complete-information extensive form of the game. For these two reasons, the most convenient way of organizing the experimental data is as empirical aggregate distributions of moves (*Pass*, *Take*), in each of the four decision nodes of the complete-information centipede game.

The use of reduced normal form strategies was fruitful for employing the framework of DFL and for defining the appropriate solution concepts. Now we need to examine what THSCE predicts about aggregate distributions of moves. The idea is the following: imagine that there are large populations of individuals, who are randomly matched, and play the centipede game as predicted by the THSCE. We randomly take the results from a large number of matches, and compile aggregate distributions of moves in each decision node. Assume further that we are unable to distinguish between different types, so we pull the data from different types together. Then, what kinds of distributions of moves could we observe?

We shall restrict attention to a specific set of strategy profiles: $\Sigma^* = \{\sigma \in \Sigma_1 \times \Sigma_2 : s_1(P_1P_3|\sigma_1) \neq 0, s_2(T_2|\sigma_2, p) \neq 1\}$. This set contains the strategy profiles such that all decision nodes of the centipede game are reached with positive probability. We define the functions $\pi_i : \Sigma^* \to \Delta(\{P_i, T_i\})$ for $i = 1, 2, 3, 4$, as follows:

1. $\pi_1(P_1|\sigma) = 1 - \sigma_1(T_1)$
2. $\pi_2(P_2|\sigma) = 1 - \xi \cdot \sigma_2(T_2|\theta_2^2) - (1-\xi) \cdot \sigma_2(T_2|\theta_2^1)$
3. $\pi_3(P_3|\sigma) = \sigma_1(P_1P_3)/[1 - \sigma_1(T_1)]$
4. $\pi_4(P_4|\sigma) = [\xi \cdot \sigma_2(P_2P_4|\theta_2^2) + (1-\xi) \cdot \sigma_2(P_2P_4|\theta_2^1)]/[1 - \xi \cdot \sigma_2(T_2|\theta_2^2) - (1-\xi) \cdot \sigma_2(T_2|\theta_2^1)]$

Again, the notation $\pi_i(P_i|\sigma)$ represents the probability that $\pi_i(\sigma)$ assigns to P_i. These functions have the following interpretation. Let the strategy profile σ be played in a given period. The number $\pi_i(P_i|\sigma)$ is the fraction of the population (to whom information set i belongs) that chooses the move P_i, conditional on reaching information set i, in this period.[25]

In particular, the conditional fraction of *Pass* in the first information set of the game is simply equal to the fraction of population 1 that chooses P_1T_3 or P_1P_3. This fraction is equal to $1 - \sigma_1(T_1)$. Similarly, in the second information set of the game, the conditional fraction of *Pass* is equal to the fraction of population 2 that does not choose T_2. In the third information set things complicate a bit. The conditional fraction of *Pass* is not the fraction of population 1 that chooses action P_1P_3. The reason is that the individual 1's who have reached the third information set do not constitute a representative sample of population 1. For, only individuals who do not choose action T_1 may have reached that node. Accordingly, the appropriate fraction is given by the proportion of individual 1's who choose action P_1P_3, among all individual 1's who choose either action P_1P_3 or action P_1T_3.[26] This is given by

[25] We shall call $\pi_i(T_i|\sigma)$ the theoretical "Conditional *Take* Fraction" in node i. This prediction will be compared with the empirical "Conditional *Take* Fraction" from our data.

[26] To put it differently, the individual 1's who reach node 3 are a representative sample of the subpopulation of individual 1's who choose either action P_1P_3 or action P_1T_3. Therefore, the conditional fraction of move P_3 must be equal to the proportion of this subpopulation that chooses action P_1P_3.

$\sigma_1(P_1P_3)/[1-\sigma_1(T_1)]$. The same logic holds for the conditional fraction of *Pass* in the fourth information set. The denominator in the right-hand side of the fourth equation is the fraction of population 2 that does not choose T_2. The numerator is the fraction of population 2 that chooses P_2P_4.

Using several propositions, we shall get restrictions on the conditional fractions of *Pass* and *Take* that can be observed in a THSCE. After introducing the experimental data, we shall then check if these restrictions hold for the empirical distributions of moves. The propositions refer both to the initial-payoff game of Figure 1, and to the modified-payoff game of Figure 2, and they consider both "information feedback" environments. Propositions $1-3$ are proven in the appendix.

The game with the initial payoffs

It will be understood that the following three propositions concern our incomplete information centipede game with the initial payoffs. Propositions 1 and 2 consider the environment without aggregate information, and Proposition 3 pertains to the environment with full or partial aggregate information.

Proposition 1: Let σ be a strategy profile such that $\sigma_1(T_1) \neq 1$. Then, σ is a type-heterogeneous SCE only if $\sigma \in \Sigma^*$, $\pi_i(T_i|\sigma) \neq 1$ for $i = 1, 2, 3, 4$, and $\pi_2(P_2|\sigma) \cdot \pi_4(P_4|\sigma) = \xi$.

Proposition 2: Let σ be a strategy profile such that $0 < \pi_i(T_i|\sigma) < 1$ for $i = 1, 2, 3, 4$. Then, σ is a type-heterogeneous SCE only if $1/7 \leq \pi_4(P_4|\sigma) \leq 7\xi$, $\pi_3(P_3|\sigma) \geq 1/7$, and $1/7 \leq \pi_2(P_2|\sigma) \leq 7\xi$.

Proposition 3: Let σ be a strategy profile such that $\sigma_1(T_1) \neq 1$. Then, σ is a type-heterogeneous SCE with aggregate information only if $\pi_1(T_1|\sigma) = 0, \pi_2(P_2|\sigma) = 7\xi, \pi_3(P_3|\sigma) = 1/7$ and $\pi_4(P_4|\sigma) = 1/7$.

Notice that, in Proposition 1, we have added the redundant condition that σ belongs to Σ^*, in order to emphasize the fact that the functions $\pi_i, i = 1, 2, 3, 4$ are defined for all the THSCE profiles described in the propositions. This indicates that restricting attention to Σ^* does not rule out any THSCE, except the trivial one where $\sigma_1(T_1) = 1$.

The game with the modified payoffs

It will be understood that the following three propositions concern the four-move centipede game with the modified payoffs. This game is similar to the game of Figure 3, except in the last terminal node of the subgame where player 2 is selfish, where payoffs are $(9, 3)$, instead of $(9.6, 2.4)$. The modification is slight, so the conditions in Propositions $1-3$ change only partially. In particular, Proposition 1 holds as it is. We also have the following new propositions that apply to the modified-payoff game:

Proposition 2*: Let σ be a strategy profile such that $0 < \pi_i(T_i|\sigma) < 1$, for $i = 1, 2, 3, 4$. Then, σ is a type-heterogeneous SCE only if $0.157 \leq \pi_4(T_4|\sigma) \leq 7\xi$, $\pi_3(P_3|\sigma) \geq 1/7$, and $1/7 \leq \pi_2(P_2|\sigma) \leq \xi/0.157$.

Proposition 3*: Let σ be a strategy profile such that $\sigma_1(T_1) \neq 1$. Then, σ is a type-heterogeneous SCE with aggregate information only if $\pi_1(T_1|\sigma) = 0, \pi_2(P_2|\sigma) = \xi/0.157$, $\pi_3(P_3|\sigma) = 1/7$ and $\pi_4(P_4|\sigma) = 0.157$.

The proofs are similar to the ones of the game with the initial payoffs, and are omitted. Evidently, in the presence of aggregate information, the set of THSCE outcomes is smaller. Notice that the model does not predict different behavior when partial, rather than full information is revealed. Equilibrium only requires correct beliefs about opponents' distribution of moves. In section 6, we shall use an estimate of the parameter ξ and a sensible refinement, and show that the conditions of the propositions describe the data relatively accurately.

4 The experiments

Twelve experimental sessions were conducted at the California Social Science Experimental Laboratory (CASSEL) at UCLA, between March and October of 2007. All subjects were UCLA students and the vast majority was undergraduate students. There were nine sessions with $n = 30$ (n is the number of participants), two sessions with $n = 28$, and one session with $n = 26$. Each subject played $\frac{n}{2}$ rounds of the four-move centipede game, plus three practice rounds. A rotating matching scheme was used, as in MP, and the subject pool was divided into two groups of $\frac{n}{2}$, the composition of which was fixed throughout each session.[27] Each participant was matched with each member of the other group exactly once. All the information about the structure of the game and the matching details was made public knowledge to subjects, since the instructions were read in public. Subjects were paid the amount that they accumulated in all real rounds, plus a \$5 participation fee, and each experimental currency unit corresponded to one US dollar. Participants did not have particular difficulties understanding the game, and also had many opportunities for learning, during the practice rounds and the (usually 15) real rounds of the game. Average payoffs were equal to \$18.73 and each session lasted for about 50 minutes.

Table 1 shows the basic features of all 12 sessions. The game played in the first seven sessions (the initial-payoff sessions), was exactly the one described in Figure 1 (the amounts refer to US dollars).[28] In sessions $NIR1$ and $NIR2$ the treatment was called "No Information Revelation" (NIR). This treatment was essentially a replication of the four-move sessions of MP. In sessions $FIR1$ and $FIR2$ the treatment was called "Full Information Revelation" (FIR). Subjects received information about how both groups played in the previous round. In particular, during each round, all subjects saw the proportions of *Pass* and *Take*, in each of the four decision nodes of the game, in the previous round. For example, during round 10, in the first decision node, all subjects saw the fraction of the members of the GREEN group that chose *Pass* or *Take*, in this particular node, during the ninth round. In the second decision node, all subjects saw the fractions of the members of the YELLOW group that chose *Pass* or *Take*, in this node, in the ninth round. Similar information was shown in all

[27] The two groups were labeled the "GREEN" group and the "YELLOW" group. The members of the GREEN group always had the role of player 1 in the centipede game, and the members of the YELLOW group always had the role of player 2. Hence, each subject played the role of only one player (1 or 2) throughout all rounds.

[28] These sessions, therefore, had the same payoff structure as in MP, but dollar payoffs were 50% higher in each terminal node, in order to account for the difference in purchasing power, due to the temporal distance between the experiments.

Table 1: Characteristics of Each Experimental Session

Session	# of Subjects	Aggregate Info	# of Matches	Payoffs
NIR1	30	NO	225	Similar with MP
NIR2	28	NO	196	Similar with MP
FIR1	30	FULL	225	Similar with MP
FIR2	30	FULL	225	Similar with MP
PIR1	30	PARTIAL	225	Similar with MP
PIR2	28	PARTIAL	196	Similar with MP
PIR3	30	PARTIAL	225	Similar with MP
NIR1-M	30	NO	225	Modified
NIR2-M	26	NO	169	Modified
FIR1-M	30	FULL	225	Modified
FIR2-M	30	FULL	225	Modified
FIR3-M	30	FULL	225	Modified

the other decision nodes.[29]

In sessions $PIR1, PIR2$ and $PIR3$, the treatment was called "Partial Information Revelation" (PIR). Subjects received the same kind of information as in treatment FIR, but only for the opposite group. For example, in round 5, all GREEN (YELLOW) subjects were shown the fractions of the YELLOW (GREEN) group that chose *Pass* or *Take*, in round 4, in both nodes where YELLOW (GREEN) moves. Subjects could not see the fractions that described the past behavior of their own group. We introduced treatment PIR in order to examine a prediction of our model, namely that only information about opponents matters, so information about the behavior of own population should be irrelevant for behavior.

In the remaining five sessions, the payoff functions were slightly modified. In particular, subjects played the game illustrated in Figure 2. There were two modified-payoff treatments, $NIR-M$ and $FIR-M$. In the two sessions of treatment $NIR-M$, there was no aggregate feedback. In the three sessions of treatment $FIR-M$, full aggregate information was revealed, with "full information" having the same meaning as above. The modified-payoff treatments were introduced in order to test the hypothesis that initial behavior serves as an equilibrium selection device. As we shall explain in sections 6 and 7, the evolution of play in the sessions with aggregate feedback seems to depend on initial conditions, and the modification in payoffs changes these conditions.

5 Experimental results

In this section, we shall present the aggregate data in a manner that helps us illustrate differences across treatments. Each experimental match will be characterized by the terminal node that it reached. Our main descriptive statistic will be the distribution (of total matches)

[29] Of course, since late decision nodes were not necessarily reached in every match, subjects saw information only about those matches that reached a given node in the previous round. If a given decision node was not reached in any match, during the previous round, no proportions were shown.

Figure 4: Distributions over Terminal Nodes, NIR, FIR and PIR, all Rounds

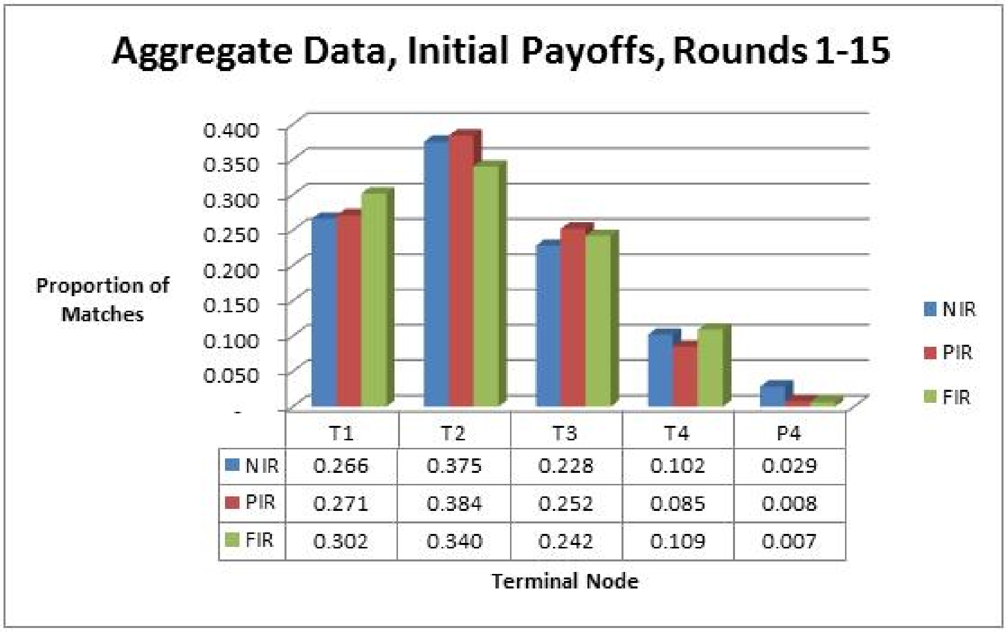

Note: In sessions $NIR2$ and $PIR2$ there were only 14 rounds. For these sessions, the data from rounds $1-14$ are used here. There are 421 matches in treatment NIR, 450 matches in treatment FIR, and 646 matches in treatment PIR.

across terminal nodes.[30] Figures $4-7$ display the aggregate data, in terms of distributions over terminal nodes. We are also interested in illustrating differences between short-run and long-run behavior. Thus, we aggregate the data over all rounds, as well as over "late" rounds only (starting from round 11). Since our main focus is on equilibrium behavior, in our analysis of the results we will mainly refer to the data from late rounds. We name each terminal node according to the last move required to reach that node (see Figure 1). So, the terminal nodes, from first to last, are denoted $T1$, $T2$, $T3$, $T4$, and $P4$.

Another measure that will interest us is the average of total monetary payoffs. In a given terminal node, total payoffs are simply the sum of the monetary earnings of the two players. Taking their average over terminal nodes (weighted by the respective frequencies), they provide a measure of average social payoffs achieved, and also of the degree to which subjects tended to pass to late nodes in the game.

Before introducing the results for each treatment separately, we should note that our data exhibit some general similarities to the results from previous experiments of the centipede game. A stylized fact from previous experiments is that the Conditional *Take* Fractions (CTF)[31] increase as we move from the first to the last decision node of the game. In our data, this was true for all treatments (see Tables $5-9$ in section 6). However, the data in our new treatments have some substantial novel features, which we shall now present.

[30] In section 6, where the focus will be on the comparison of the predictive performance of the various theories, we will present the same data as distributions of moves in each decision node.

[31] The Conditional *Take* Fraction, in a particular decision node, is the number of matches where *Take* was chosen in this node, divided by the number of all matches that reached the node.

Figure 5: Distributions over Terminal Nodes, NIR, FIR and PIR, Rounds 11-15

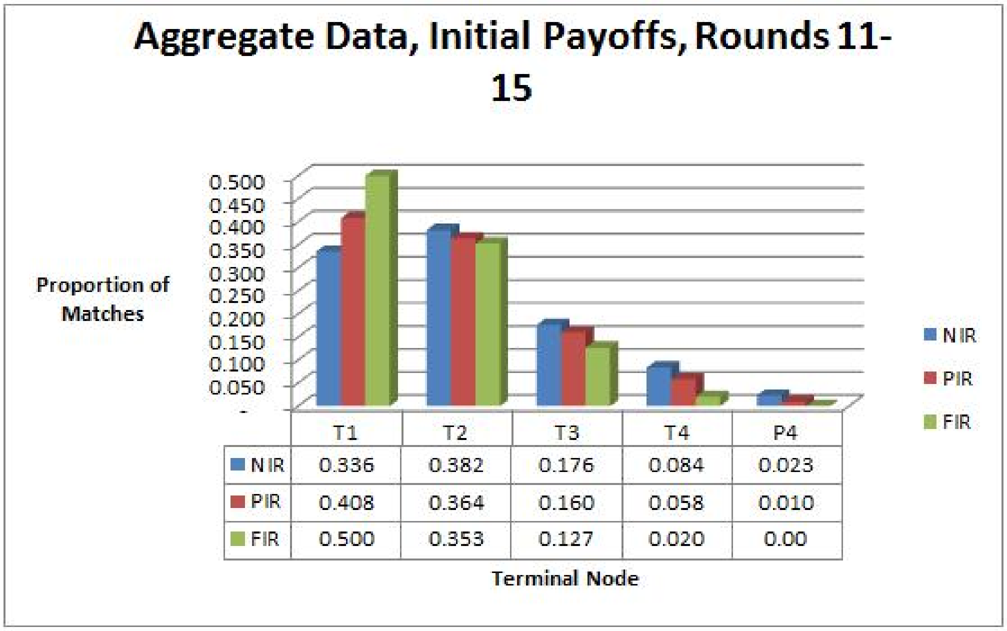

Note: In sessions $NIR2$ and $PIR2$ there were only 14 rounds. For these sessions, the data from rounds $11-14$ are used here. There are 131 matches in treatment NIR, 150 matches in treatment FIR, and 206 matches in treatment PIR.

5.1 Treatments NIR and FIR

Figures 4 and 5 illustrate the distributions over terminal nodes of treatments NIR and FIR. Figure 4 shows that there are small differences between treatments NIR and FIR, if all rounds are considered. Figure 5, which considers only late rounds, reveals larger differences. For example, the percentage of total matches that ended in terminal node $T1$ (which corresponds to the "bad" THSCE outcome of our model, as well as to the unique Nash equilibrium outcome of the game with selfish payoffs) is about 50% in FIR and about 33% in NIR. Convergence to terminal node $T1$ is very strong in late rounds, much stronger than in MP.

Tables 2 and 3 contain the main statistical tests, and we shall frequently refer to them. To make statistical tests, we assumed that each match is independent of the others.[32] Table 2 presents tests, which use data in late rounds only.[33] We see that the higher frequency

[32]Since each subject is coupled with each member of the other group exactly once, each match is unique. If we drop the assumption that matches in different rounds are independent, we can only run tests for a single round. For completeness, and because the choice of any single round would seem arbitrary, we also performed tests $1-6$ (from Table 1) individually for each late round. We used Fisher's exact test, and we report the p-values in Table 4, in page 27. The other tests cannot be performed for individual rounds. Sample sizes in each round are too small to use chi-square tests, and the frequency of terminal node $P4$ is also too low in each round to run reasonable tests.

[33]For example, test 1 examines the null hypothesis that the fraction of total matches that end in the first terminal node is the same for treatments NIR and FIR. Test 7 examines the null hypothesis that the distributions across terminal nodes do not differ in treatments NIR and FIR.

of terminal node $T1$ in treatment FIR, relative to NIR, is statistically significant (test 1). Moreover, the whole distributions, in these two treatments, differ substantially, and the chi-square test shows that this difference is significant (test 7). Furthermore, average total payoffs per match are \$2.13 in treatment NIR, and \$1.4 in treatment FIR, (always in rounds $11-15$) (t-test, two-tailed, $p = 0.0004$). Therefore, full aggregate information seems to have a negative effect on social payoffs. To sum up, the evidence indicates that subjects' behavior is different when full information is provided.

Furthermore, in the two sessions of our control treatment, subjects seem to behave differently than in the experiments of MP. The percentage of terminal node $T1$, in rounds $6-10$, was 29% and 38.5%, in sessions $NIR1$ and $NIR2$, respectively. In MP's four-move treatment, the analogous percentage was 6% (in two sessions) and 10% (in one session). Pooling the data from different sessions together, the χ^2 test rejects the null hypothesis of homogeneity of the distributions in our NIR treatment and in MP's treatment (test 8).

It is also noteworthy that in treatment FIR, very few matches reached the last terminal node ($P4$). This node is particularly interesting, because it involves a dominated choice. However, the frequency of this terminal node is so low, that data from all real rounds need to be pooled together, in order to perform statistical tests. These tests are reported in Table 3. A simple test of differences in proportions finds that a significantly higher fraction of total matches reaches terminal node $P4$ in NIR relative to FIR (test 17). We will also perform Fisher's exact test[34] whenever the expected frequency for any category is very low, and the contingency table is 2×2.[35] This test also finds a significant difference in the proportions of $P4$ (test 16). In addition, the CTF in the last (fourth) decision node is 0.942 in treatment FIR and 0.782 in treatment NIR. The χ^2 test indicates that the difference is statistically significant (test 21).

5.2 Treatment PIR

Figures 4 and 5 also display the distribution over terminal nodes of treatment PIR. We will mainly consider behavior in late rounds. As in treatment FIR, convergence towards the first terminal node was observed, but only in sessions $PIR2$ and $PIR3$. In session $PIR1$, there were signs of convergence to the high-payoff THSCE outcome (we shall return to this in section 6). As a result, the fraction of node $T1$ is not statistically different in PIR from the respective fraction in either NIR or in FIR (tests 2 and 3). In general, the hypothesis of homogeneity of the distributions in NIR, FIR and PIR can only be rejected at the 10% level (test 9).[36] Moreover, average total payoffs per match in PIR are equal to \$1.79, which is higher than in FIR (t-test, two tailed, $p = 0.0055$) and somewhat lower than in NIR (t-test, two-tailed, $p = 0.129$).

[34] We made the calculations using the software of Preacher and Briggs (2001).

[35] A well-known weakness of z and chi-square testing is its inappropriateness when some category has very low "expected frequency", as is the case in our data. Conventional wisdom in the statistics literature says that Fisher's exact test is more appropriate for small samples, and chi-square tests for large samples. See D' Agostino, Chase, and Belanger (1988) and Sahai and Khurshid (1995) for an excellent review of the appropriate methods.

[36] The pairwise tests give similar results. The distribution in NIR is not statistically different from the distribution in PIR, and the distribution in FIR is also not statistically different from the distribution in PIR (tests 10 and 11).

Table 2: Statistical Tests Using Data Aggregated over Rounds $11-15$

Test	Object	Treatments	Test Statistic	P-Value
1	Π1	NIR(0.33) and FIR(0.5)	z=2.77	0.005
2	Π1	NIR(0.33) and PIR(0.4)	z=1.32	0.18
3	Π1	FIR(0.5) and PIR(0.4)	z=1.55	0.12
4	Π1	NIR(0.33) and NIR-M(0.096)	z=20	< 0.001
5	Π1	NIR-M(0.096) and FIR-M(0.18)	z=2.06	0.038
6	Π1	FIR(0.5) and FIR-M(0.18)	$\chi^2 = 42.5$	< 0.001
7	F	NIR and FIR	$\chi^2 = 14.89$	0.004
8	F	NIR and NIR_{MP}(rounds 6-10)	$\chi^2 = 35.8$	< 0.001
9	F	NIR, FIR and PIR	$\chi^2 = 15.02$	0.059
10	F	NIR and PIR	$\chi^2 = 2.98$	0.56
11	F	FIR and PIR	$\chi^2 = 6.81$	0.14
12	F	NIR and NIR-M	$\chi^2 = 26.47$	< 0.001
13	F	NIR-M and FIR-M	$\chi^2 = 14.51$	0.006
14	F	PIR1 and PIR(2+3)	$\chi^2 = 75.1$	< 0.001
15	F	FIR2-M and FIR(1+3)-M	$\chi^2 = 90.1$	< 0.001

Note: "Object" refers to the parameter which is equal across treatments under the null hypothesis. "F" denotes the whole distribution over the five terminal nodes. Π1 denotes the probability that a match finishes in the first terminal node. The numbers in the parentheses are the empirical realizations. $PIR(2+3)$ denotes the pooled data from sessions $PIR2$ and $PIR3$. $FIR(1+3)-M$ has a similar meaning.

We now consider behavior in the last decision node, in rounds $1-15$. The CTF in decision node 4 in PIR is equal to 0.916. This is very similar to the CTF in FIR (test 24), and higher than the CTF in NIR (although tests 22 and 23 provide mixed results at the 5% level). Homogeneity of the CTF across all treatments NIR, PIR and FIR is rejected at the 5% level (test 25). Additionally, the proportion of total matches that ended in terminal node $P4$ in treatment NIR is significantly higher than in treatment PIR (test 19).

5.3 Treatments NIR-M and FIR-M

Recall that in the modified-payoff treatments, YELLOW subjects who choose *Pass* in their last decision node have somewhat higher monetary payoffs than before ($3 instead of $2.4). To examine the effects of aggregate information in this setting, we will compare behavior in the sessions with the new payoffs and no aggregate information ($NIR-M$), to behavior in the sessions with the new payoffs and full information revelation ($FIR-M$). We will also discuss the important differences in subjects' behavior for a fixed level of information feedback (NIR vs. $NIR-M$ and FIR vs. $FIR-M$).

Figures 6 and 7 display the distributions over terminal nodes of treatments $NIR-M$ and $FIR-M$. A comparison of the distributions in treatments NIR and $NIR-M$ shows an important independent effect of the slight change in payoffs (test 12). Moreover, a very low

Table 3: Statistical Tests Using Data Aggregated over all Rounds $(1-15)$

Test	Object	Treatments	Test Statistic	P-Value
16	Π5	NIR(0.028) and FIR(0.0067)	Fisher	0.017
17	Π5	NIR(0.028) and FIR(0.0067)	z=2.47	0.013
18	Π5	NIR(0.028) and NIR-M(0.045)	$\chi^2 = 1.69$	0.19
19	Π5	NIR(0.028) and PIR(0.007)	Fisher	0.011
20	Π5	NIR-M(0.045) and FIR-M(0.091)	$\chi^2 = 7.65$	0.006
21	$\Pi5/[\Pi5+\Pi4]$	NIR(0.782) and FIR(0.942)	$\chi^2 = 5.7$	0.017
22	$\Pi5/[\Pi5+\Pi4]$	NIR(0.782) and PIR(0.916)	$\chi^2 = 4.1$	0.042
23	$\Pi5/[\Pi5+\Pi4]$	NIR(0.782) and PIR(0.916)	Fisher	0.064
24	$\Pi5/[\Pi5+\Pi4]$	FIR(0.942) and PIR(0.916)	$\chi^2 = 0.27$	0.6
25	$\Pi5/[\Pi5+\Pi4]$	NIR, FIR and PIR	$\chi^2 = 7.7$	0.021
26	$\Pi5/[\Pi5+\Pi4]$	NIR(0.782) and NIR-M(0.84)	$\chi^2 = 0.923$	0.336
27	$\Pi5/[\Pi5+\Pi4]$	NIR-M(0.84) and FIR-M(0.69)	$\chi^2 = 8.7$	0.003

Note: Π4 and Π5 denote the probability that a match finishes in the fourth and fifth terminal node, respectively. The empirical values are contained in the parentheses. The number of matches that reached the last decision node is 55 for treatment NIR, 52 for treatment FIR, 60 for treatment PIR, 114 for treatment $NIR-M$, and 201 for treatment $FIR-M$.

fraction of total matches ends in terminal node $T1$ in treatment $NIR-M$. This fraction significantly differs from treatment NIR (test 4). Importantly, the effects of aggregate information on subjects' behavior now seem markedly different. Average total payoffs in $FIR-M$ and $NIR-M$ are \$3.37 and \$2.92 respectively (t-test, one-tailed, $p = 0.061$). Hence, in the modified-payoff setting, full information tends to somewhat increase average payoffs, rather than decrease them.[37] Furthermore, average total payoffs per match in NIR and in $NIR-M$ differ significantly (t-test, one-tailed, $p = 0.002$).

In addition, aggregate information does not cause convergence to terminal node $T1$ in sessions $FIR1-M$ and $FIR3-M$, but only in session $FIR2-M$ (see section 6). This large difference between sessions, hidden in the pooled data, causes a paradoxical combination of higher frequency for both terminal node $T1$ and terminal node $P4$, when aggregate feedback is provided. More games end in terminal node $T1$ in $FIR-M$ relative to $NIR-M$, a difference which is statistically significant at the 5% level (test 5). On the other hand, in treatment $FIR-M$, many subjects achieved very high payoffs, reaching the last or the penultimate terminal node. A significantly higher fraction of matches reached terminal node $P4$ in treatment $FIR-M$ compared to $NIR-M$ (test 20). Additionally, the CTF in decision node 4 in treatment $NIR-M$ is higher than in treatment NIR (0.84 vs. 0.782), but the difference is not statistically significant (test 26). The CTF in $FIR-M$ is only 0.69, significantly different than in $NIR-M$ (test 27).

[37] Recall that average total payoffs in treatment FIR were equal to \$1.4. In treatment $FIR-M$, which differs from treatment FIR in a minor way, average total payoffs are more than double (t-test, one-tailed, $p < 0.001$).

Figure 6: Distributions over Terminal Nodes, $NIR - M$ and $FIR - M$, all Rounds

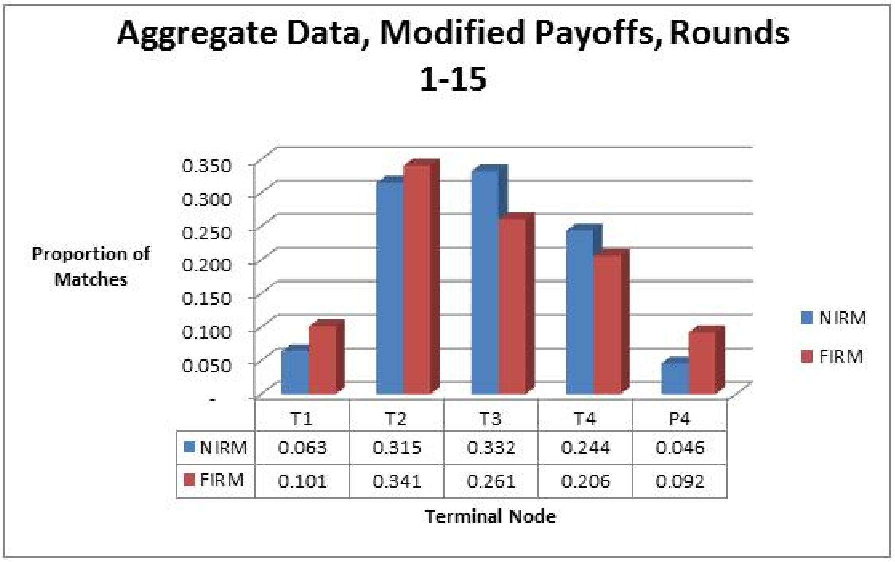

Note: In session $NIR2 - M$, there were 13 rounds. For this session, the data from rounds 1-13 are used here. There are 394 matches in treatment $NIR - M$ and 675 matches in treatment $FIR - M$.

6 The performance of theoretical predictions

We shall examine whether the predictions of our model, and the predictions of three different, frequently used solution concepts, are consistent with the qualitative aspects of the data. In particular, we shall compare the empirical CTF, in each decision node, with the theoretical values of the CTF.[38] We will consider the Nash equilibrium prediction for the two games of Figures 1 and 2, and the AQRE prediction for the same games. Moreover, we shall consider the sequential equilibrium (SE) prediction for the incomplete information game of Figure 3, and also for the same incomplete information game with the modified payoffs.

We should emphasize that these three solution concepts predict that aggregate information should have no effect on equilibrium behavior. In all of these approaches, the behavior of the other population is assumed to be correctly anticipated in equilibrium. Feedback and learning play no explicit role in forming agents' beliefs. Thus, the predictions of these concepts depend only on the extensive-form game. On the contrary, our model has different predictions for the two different feedback environments.

Tables $5 - 9$ juxtapose the observed distributions of moves from our data (in rounds $11 - 15$), in each of the four decision nodes of our game, with the theoretical values of these distributions. Each table corresponds to a different treatment. By ϕ_i, where $i = 1, 2, 3, 4$, we denote the CTF in node i. Since our model has multiple equilibria, the predictions are

[38] To be more precise, the theoretical CTF, according to some model, is the fraction of *Take*, that one would observe by looking at the behavior of large populations of agents, who individually play according to the equilibrium prediction. Consider, for example, a standard incomplete information model. For a given decision node, the theoretical CTF is given by the average (over types) of the equilibrium probabilities of *Take*, in this decision node, weighted by each type's probability of reaching this node (see footnote 42). Recall that for our model the theoretical CTF in node i is equal to $\pi_i(T_i|\sigma)$.

Figure 7: Distributions over Terminal Nodes, $NIR-M$ and $FIR-M$, Rounds 11-15

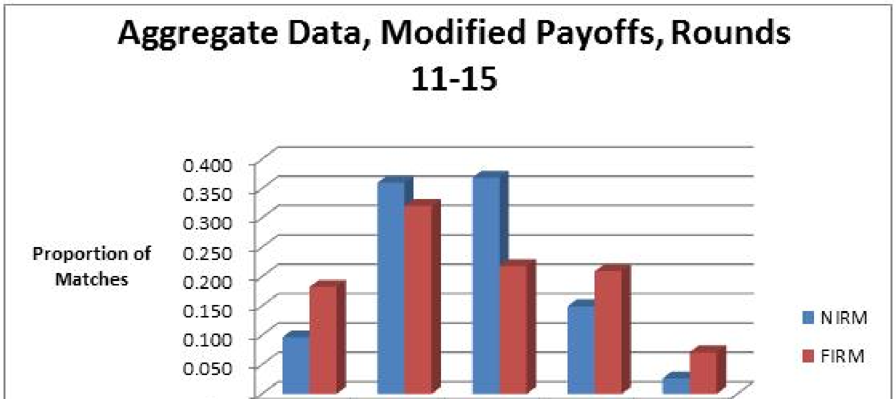

	T1	T2	T3	T4	P4
NIRM	0.096	0.360	0.368	0.149	0.026
FIRM	0.182	0.320	0.218	0.209	0.071

Note: In session $NIR2-M$, there were 13 rounds. For this session, the data from rounds 11-13 are used here. There are 114 matches in treatment $NIR-M$ and 225 matches in treatment $FIR-M$.

often presented as inequalities.

6.1 Nash equilibrium, AQRE and SE

In the Nash equilibrium of the centipede game, the CTF in decision node 1 is equal to one. To calculate the AQRE, we used the program Gambit (McKelvey, McLennan, and Turocy, 2007) and then estimated the main parameter λ with maximum likelihood. One value of λ for each of the extensive forms of Figures 1 and 2 was estimated. The estimated values of λ were 0 and 1.197, respectively. Using these values, Gambit yields the theoretical CTFs shown in Tables $5-9$.

In order to get exact predictions for our model, as well as for the SE model, we need to estimate the fraction of altruists ξ. To do this, we apply the equation $\pi_2(P_2|\sigma)\cdot\pi_4(P_4|\sigma)=\xi$, from Proposition 1, to the data from the four-move sessions of MP (in rounds $6-10$, which are the last 5 rounds). Since $\pi_i(T_i|\sigma)$ is the theoretical CTF in node i, the empirical estimate of the fraction of altruists would be $\hat{\xi}=(1-\phi_2)(1-\phi_4)=(0.51)(0.18)\simeq 0.09$. This value is not very different from 0.05, which was estimated in the model of MP.

Standard calculations show that the incomplete information game of Figure 3 has a unique SE, with the following behavior strategies for player 1 and selfish player 2: $\sigma_1 = [(0,1);(\frac{6}{7},\frac{1}{7})]$, $\sigma_2^{self.} = [(\frac{1-7\xi}{1-\xi},\frac{6\xi}{1-\xi});(1,0)]$.[39] Also consider the incomplete information game with the modified payoffs.[40] For this game, there is also a unique SE, and the behavior

[39]The first (second) parenthesis in each brackets refers to the first (second) decision node of each player, and the first number in each parenthesis corresponds to the probability of *Take*.

[40]Recall that this game is as shown in Figure 3, but with payoffs $(9,3)$ instead of $(9.6,2.4)$ in the last terminal node of the selfish type.

Table 4: P-Values of Tests 1 − 6 for Individual Late Rounds

	Test 1	**Test 2**	**Test 3**	**Test 4**	**Test 5**	**Test 6**
Round 11	0.143*	0.629	0.04*	0.02	0.467	1
Round 12	0.115	0.148	0.816	0.041	0.467	< 0.001
Round 13	0.067	0.586	0.004	0.103	0.732	< 0.001
Round 14	0.008	0.265	0.103	0.302	0.712	< 0.001
Round 15	0.057	0.203	0.604	0.652	0.712	0.001

Note: In the tests marked with an asterisk, the difference between treatments is on the opposite direction than the one exhibited in the aggregated data over rounds 11 − 15, which is shown in Table 2.

Table 5: Data and Theoretical Predictions, Treatment NIR

Node	**Fraction**	**Data**	**THSCE**	**Nash**	**AQRE**	**SE**
1	$1-\phi_1$	0.665		0	0.5	1
	ϕ_1	0.335	$0 < \phi_1 < 1$	1	0.5	0
2	$1-\phi_2$	0.425	$0.143 \leq (1-\phi_2) \leq 0.63$	-	0.5	0.63
	ϕ_2	0.575	$\& \ (1-\phi_2) = 0.09/(1-\phi_4)$	-	0.5	0.37
3	$1-\phi_3$	0.378		-	0.5	0.143
	ϕ_3	0.622	$\phi_3 \leq 0.857$	-	0.5	0.857
4	$1-\phi_4$	0.214	$0.143 \leq (1-\phi_4) \leq 0.63$	-	0.5	0.143
	ϕ_4	0.786		-	0.5	0.857

strategies are: $\sigma_1 = [(0,1);(\frac{6}{7},\frac{1}{7})]$, $\sigma_2^{self.} = [(\frac{1-6.5\xi}{1-\xi},\frac{5.5\xi}{1-\xi});(1,0)]$. Using the value $\hat{\xi} = 0.09$, we calculate the CTFs induced by the SE strategies.[41]

The numerical predictions from these three equilibrium models, shown in Tables 5 − 9, fail to capture the treatment effect of aggregate feedback. In addition, even in the treatments with no aggregate feedback, the Nash equilibrium prediction performs very poorly (Tables 5 and 8). Moreover, in the treatments with the initial payoffs (Tables 5 − 7), the AQRE predicts fifty-fifty play in all nodes, whereas the SE performs slightly better. Overall, these

[41] These CTFs can be found as follows. In decision nodes 1 and 3, since player 1 has only one possible type, the distributions of moves simply coincide with the behavior strategies of this player. Regarding nodes 2 and 4, recall that all altruists choose *Pass*. The relative proportions of altruistic and selfish individual 2's, who move in node 2, are ξ and $(1-\xi)$. So, the CTF, which is induced by the SE, is simply $(1-\xi)\cdot[(1-7\xi)/(1-\xi)] = 1-7\xi$ (which is 0.37, substituting the estimated value of ξ). Thus, a fraction ξ of the total population of 2's reaches node 3 and is altruistic, and a fraction 6ξ of the total population of 2's reaches node 3 and is selfish. Some proportion of this large number of individual 2's, who have reached node 3, is randomly chosen (their opponents choose *Pass* in node 3), and they get to play in node 4. Accordingly, the relative proportion of altruists, out of all individual 2's who move in node 4, is still 1/7. Since all selfish individuals choose *Take* and all altruists choose *Pass* in node 4, the predicted fraction of *Take* is 6/7. The same reasoning holds for the incomplete information game with the modified payoffs.

Table 6: Data and Theoretical Predictions, Treatment FIR

Node	Fraction	Data	THSCE	Nash	AQRE	SE
1	$1-\phi_1$	0.5	0	0	0.5	1
	ϕ_1	0.5	1	1	0.5	0
2	$1-\phi_2$	0.3	-	-	0.5	0.63
	ϕ_2	0.7	-	-	0.5	0.37
3	$1-\phi_3$	0.137	-	-	0.5	0.143
	ϕ_3	0.863	-	-	0.5	0.857
4	$1-\phi_4$	0	-	-	0.5	0.143
	ϕ_4	1	-	-	0.5	0.857

Table 7: Data and Theoretical Predictions, Treatment PIR

Node	Fraction	Data	2+3	1	THSCE	Nash	AQRE	SE
1	$1-\phi_1$	0.593	0.405	0.92	0	0	0.5	1
	ϕ_1	0.407	0.595	0.08	1	1	0.5	0
2	$1-\phi_2$	0.385	0.17	0.551	-	-	0.5	0.63
	ϕ_2	0.615	0.83	0.449	-	-	0.5	0.37
3	$1-\phi_3$	0.398	0	0.268	-	-	0.5	0.143
	ϕ_3	0.702	1	0.632	-	-	0.5	0.857
4	$1-\phi_4$	0.143	-	0.143	-	-	0.5	0.143
	ϕ_4	0.857	-	0.857	-	-	0.5	0.857

Note: "2+3" denotes sessions $PIR2$ and $PIR3$ (pooled) and "1" denotes session $PIR1$. "Data" still denotes all the three sessions pooled.

models seem inappropriate for capturing several important features of the data.

6.2 Our model

Now, consider the predictions of our model, using the estimated value $\hat{\xi} = 0.09$. In our treatments without aggregate information, all terminal nodes are reached. Propositions 1 and 2 put restrictions on the CTFs that should be observed in a strategy profile where all terminal nodes are reached, if behavior is compatible with THSCE. As we see in Table 5, these restrictions are met by our data in treatment NIR. Regarding treatment $NIR - M$, the predictions of Proposition 2*, shown in Table 8, are also generally accurate. Remarkably, the condition $(1 - \phi_2)(1 - \phi_4) = 0.09$ holds almost exactly for both NIR and $NIR - M$.

In the environment with aggregate information, for both payoff functions, there are two equilibrium outcomes. The "good" equilibrium outcome is described by the conditions in Proposition 3 or 3*, and the "bad" equilibrium outcome is the outcome where $\sigma_1(T_1) = 1$. We shall use the bad outcome as a prediction for treatments FIR and PIR, and the good

Table 8: Data and Theoretical Predictions, Treatment $NIR - M$

Node	Fraction	Data	THSCE	Nash	AQRE	SE
1	$1 - \phi_1$	0.903		0	0.803	1
	ϕ_1	0.097	$0 < \phi_1 < 1$	1	0.196	0
2	$1 - \phi_2$	0.601	$0.143 \leq (1 - \phi_2) \leq 0.573$	-	0.756	0.585
	ϕ_2	0.398	$\& \ (1 - \phi_2) = 0.09/(1 - \phi_4)$	-	0.244	0.415
3	$1 - \phi_3$	0.323		-	0.385	0.143
	ϕ_3	0.677	$\phi_3 \leq 0.857$	-	0.615	0.857
4	$1 - \phi_4$	0.15	$0.157 \leq (1 - \phi_4) \leq 0.63$	-	0.104	0.143
	ϕ_4	0.85		-	0.896	0.857

Table 9: Data and Theoretical Predictions, Treatment $FIR - M$

Node	Fraction	Data	1+3	2	THSCE	Nash	AQRE	SE
1	$1 - \phi_1$	0.818	0.967	0.52	1	0	0.803	1
	ϕ_1	0.182	0.033	0.48	0	1	0.196	0
2	$1 - \phi_2$	0.608	0.71	0.231	0.573	-	0.756	0.585
	ϕ_2	0.392	0.29	0.769	0.427	-	0.244	0.415
3	$1 - \phi_3$	0.56	0.583	0.333	0.143	-	0.385	0.143
	ϕ_3	0.44	0.417	0.667	0.857	-	0.615	0.857
4	$1 - \phi_4$	0.254	0.267	0	0.157	-	0.104	0.143
	ϕ_4	0.746	0.733	1	0.843	-	0.896	0.857

Note: "1+3" denotes sessions $FIR1{-}M$ and $FIR3{-}M$ (pooled) and "2" denotes session $FIR2{-}M$. "Data" still denotes all the three sessions pooled.

outcome (described in Proposition 3*) as a prediction for treatment $FIR - M$. The reason is that, as we shall argue, initial behavior serves as an equilibrium selection device, and behavior in initial rounds was very different in $FIR - M$ than initial behavior in the other two treatments. This selection device is supported by evidence discussed in section 7, which shows that aggregate information induces strong path dependence. This may lead sessions, which start with small differences, to end up with very different long-run behavior.[42]

We argue that initial differences in the CTF in decision node 4 might be critical for long-run play. In particular, if this fraction in early rounds is less than 6/7, initial conditions should be viewed as "optimistic".[43] Table 10 illustrates the initial fractions of *Pass* in decision

[42] This was the motivation for introducing the modified-payoff treatments. We hypothesized that the small modification in payoffs could change initial behavior in an "optimistic" way, and it could lead to the high-payoff equilibrium.

[43] To illustrate the argument, consider the decision of a GREEN subject in decision node 3. If he expects a fraction of *Pass* lower than 1/7 in decision node 4, his optimal choice in node 3 is *Take*. When aggregate feedback is provided, if the CTF in node 4 has exceeded 6/7 in the previous round, the agent will tend

node 4, for the all the sessions with aggregate feedback. The initial CTF is generally much higher in the sessions of PIR and FIR, than in the sessions of $FIR - M$. We therefore argue that initial conditions were favorable for reaching the good equilibrium outcome in treatment $FIR - M$, but not in treatment FIR or PIR.

The THSCE prediction for treatments FIR and PIR, illustrated in Tables 6 and 7, is generally consistent with the observed convergence to early terminal nodes. However, convergence to terminal node $T1$ is not complete. Furthermore, Table 8 shows that for treatment $FIR-M$ the predictions are reasonably accurate, but only for the first two decision nodes. Hence, the model does not seem to perform very well. However, the pooled data of treatments PIR and $FIR - M$ might be misleading, because different individual sessions seem to have converged to different equilibria. If we separate the data we get a different picture (Tables 7 and 9). Play in sessions $PIR2$ and $PIR3$ exhibits strong convergence to the bad equilibrium. However, session $PIR1$ seems to converge very close to the good equilibrium of the initial-payoff game (as can be verified, by plugging the value $\xi = 0.09$ in the conditions of Proposition 3). Similarly, in session $FIR2 - M$, there are strong signs of convergence to the bad equilibrium outcome, while in sessions $FIR1 - M$ and $FIR2 - M$, this is not the case. Nevertheless, even excluding session $FIR2-M$, the THSCE predictions do not seem accurate for decision nodes 3 and 4. One reason might be that repeated-game aspects may have played a role in treatment $FIR - M$ (see next session).

Our objective is not to fit the data, since a single experiment does not provide conclusive evidence. We only tried to show how a very simple deterministic model can incorporate ex-post feedback in the analysis of equilibrium behavior, which is important for analyzing experimental games.

Table 10: Initial Fractions of *Pass* in Decision Node 4, Treatments FIR, PIR, and $FIR - M$

Session	Round 1	Rounds 1-5	Session	Round 1	Rounds 1-5
$PIR1$	1/5	1/16	$FIR1 - M$	3/9	10/23
$PIR2$	1/4	1/8	$FIR2 - M$	1/2	2/12
$PIR3$	0/1	0/5	$FIR3 - M$	3/4	16/39
$FIR1$	1/5	2/27	**FIR**	**1/9**	**2/33**
$FIR2$	0/4	0/6	**PIR**	**2/10**	**2/29**
			FIR-M	**7/15**	**28/74**

Note: the numbers in bold letters are the totals for each treatment.

to choose *Take* in node 3, in the current round. So, GREEN subjects shall tend to respond to the initial behavior of the YELLOW with *Take* in node 3. The same reasoning can now be used regarding the decision of YELLOW subjects in decision node 2, who start to observe a high CTF in node 3. Continuing this backward reasoning, we see that aggregate behavior is likely to converge to the bad equilibrium outcome. On the other hand, if the initial CTF in node 4 was lower than 6/7, GREEN subjects would be inclined to choose *Pass* in node 3, and behavior need not collapse.

7 Dynamic aspects and path dependence

The large shift in the effects of aggregate feedback, observed in the modified-payoff treatments, may be due to repeated-game effects, and to the strong path dependence, induced by aggregate feedback. Repeated-game effects might explain why so many subject 2's chose *Pass* in their last decision node, in treatment $FIR - M$. Path dependence explains why small changes in initial behavior become large differences in late rounds. It also accounts for the large behavioral discrepancies found in different sessions, within a given treatment.

So far we have implicitly assumed that subjects ignore dynamic effects, and we shall argue that this is generally justified on the basis of actual behavior. In the treatments with aggregate feedback, a subject could sacrifice payoffs now, in order to affect the feedback in the next round. However, in late rounds of treatments FIR and PIR, the frequency of the choice *Pass* and average payoffs were so low, that is it difficult to imagine that subjects' behavior was forward-looking, trying to induce future cooperation. The evidence from these treatments, therefore, indicates that subjects did not seek to, or were not able to, utilize the repeated-game aspects of the interaction. We have no reason to believe that the sessions with the modified payoffs were somehow more conducive to forward-looking behavior. However, we cannot rule out the possibility that this type of behavior may have played a role in treatment $FIR - M$, given the high proportion of *Pass* in the fourth decision node. In fact, this behavior is a possible reason for the poor performance of our model in treatment $FIR - M$.

7.1 Regression analysis

Aggregate information also seems to have increased path dependence. This is indicated by the sizeable differences in behavior exhibited in different sessions of a given treatment (in particular, in treatments PIR and $FIR - M$). These differences are highly statistically significant (tests 14 and 15).

To further examine whether aggregate information induced path dependence, we run a few simple regressions. Notice that we do not observe individual subjects' strategies, so we use aggregate variables. In our first specification, we consider how the CTF in node 1, in a given round, is affected by the opponents' CTF in node 2, in the previous round. So, our first specification is: $\phi_{1,t} = b_0 + b_1\phi_{2,t-1} + \epsilon$, where $\phi_{i,t}$ is the CTF in decision node i in round t.[44] So, our dependent variable captures current "own-group" behavior, and the explanatory variable captures "other-group" behavior in the previous round. We predict a positive relationship, because subject 1's may expect that if the opponents' CTF in the previous round was high, this implies that the opponents' CTF in the current round will also be high. Clearly, the relationship should be stronger in the treatments with aggregate feedback, since the opponents' CTF in the previous round is publicly revealed in these treatments.

[44] For all our specifications, sessions of the same treatment are pooled together, and all real rounds are used. An observation consists of the CTF in node 1, for a given round t, plus the CTF in node 2, for round $t - 1$. Therefore, the CTF in node 1, for the first round of each session, is not included in any observation. Similarly, the CTF in node 2, for the last round of each session, is not included in any observation.

Our second specification is: $\phi_{1,t} = c_0 + c_1\phi_{2,t-1} + c_2\phi_{3,t-1} + \epsilon$. We include $\phi_{3,t-1}$ because sophisticated individual 1's may prefer to choose T_1 in the current round, if the CTF in decision node 3 was high in the previous round. For, they may expect that opponents will respond to this with a high fraction of *Take* in node 2 (especially when aggregate feedback is provided). Our third specification examines the relationship between the behavior of subject 1's in their last decision node (decision node 3), and the behavior of subject 2's in their last decision node (decision node 4), in the previous round: $\phi_{3,t} = \beta_0 + \beta_1\phi_{4,t-1} + \epsilon$.[45]

The results, shown in Table 11, support the idea that aggregate information reinforces path dependence. In particular, in the treatments with aggregate feedback, subjects' behavior in the current round depends strongly on opponents' behavior in the previous round (in the most relevant node). The regression coefficient of opponents' behavior in round $t-1$ is much higher in the aggregate information treatments than in the control treatments, and they are statistically significant. It should be noted that the weak results in treatment $FIR-M$ may be due to the lack of sufficient variability in the data, since in most rounds of this treatment, the CTF in decision node 1 was equal to zero.

Table 11: Regression Results

	NIR	PIR	FIR	$NIR-M$	$FIR-M$
b_0	0.073	−0.221∗	−0.19∗	−0.032	−0.108∗
b_1	0.391∗	0.904∗	0.904∗	0.278∗	0.515∗
c_0	0.02	−0.315∗	−0.177	−0.068∗	−0.082∗
c_1	0.41∗	0.698∗	0.866∗	0.18∗	0.445∗
c_2	0.070	0.242∗	0.0006	0.128∗	−0.014
β_0	-	-	-	0.54∗	0.10
β_1	-	-	-	0.021	0.565∗

* Significant at the 5% level.

In order to directly examine our claim that aggregate feedback increased path dependence, we tested whether aggregate information causes a "structural shock" in our model. We pulled all the observations of the initial-payoff treatments together, and run a Chow test of structural change in the coefficients, between the control treatment and the aggregate feedback treatments. For the first specification, the p-value of the Chow test was equal to 0.0062, and hence we reject the null hypothesis that coefficients did not change. Consequently, aggregate information seems to have increased path dependence in the initial-payoff treatments. For the second specification, the p-value of the Chow test was 0.089, hence we cannot reject the null hypothesis of equality of the coefficients. We also pulled together all the observations of the modified-payoff treatments, and run the same Chow test. We got weaker results (p-values equal to 0.232 and 0.13 for the first and the second specification, respectively). However, as we explained before, there is a lack of sufficient heterogeneity in that data. For our third specification, the p-value of the Chow test (using data from

[45]Unfortunately, the sparse availability of the data does not allow us to run this specification for all treatments. Decision nodes 3 and 4 were very rarely reached in treatments NIR, PIR and FIR.

only modified-payoff sessions, of course) was 0.0598. Hence, the hypothesis of no structural change cannot be rejected at the 5% level.

With strong path dependence, the evolution of play in sessions FIR and $FIR - M$ can be better understood. As we have shown, the payoff modification changed initial behavior, encouraging more initial passing. Once behavior started "optimistically", path dependence reinforced this tendency. Moreover, subjects were enjoying high payoffs, which may have facilitated understanding the fact that forward-looking behavior is profitable. So, it is possible that they chose to pass more than what would be justified by myopic beliefs alone.

8 Conclusions

Motivated by the literature on learning in games, we conducted an experimental investigation of the centipede game with anonymous matching. We introduced some new treatments, where subjects received information about the proportion of each group that chose each action in the previous round. We were interested in examining the equilibrium effects of aggregate information. We developed a model based on the framework of DFL, assuming a small fraction of pure altruists.

In the first set of our experiments, we used a payoff function similar to the one used in MP. We found that both full and partial information release causes convergence to the first terminal node, where social payoffs are very low. Our model captures the data relatively well. In a second set of experiments, we tested whether a change in initial conditions can alter the effects of aggregate information. In order to achieve this change in initial behavior, we slightly increased the payoffs from choosing *Pass* in node 4. In two out of three sessions with the modified payoffs, aggregate information resulted in higher average payoffs. This shift is explained by strong path dependence, which tends to magnify initial differences in behavior.

In terms of policy, the experimental results indicate that it is possible to use information release as a means of inducing more trust and achieving higher social payoffs. However, the results also raise the concern that the effectiveness of such a policy is not always guaranteed. Information seems to enhance social welfare, but only for the sessions that start with high cooperation. A policy lesson would be that only optimistic information should be revealed, in the sense that the revealed behavior should already exhibit a certain level of trust.

Regarding future study, we believe that there is more scope in examining how people respond to aggregate information. For example, revealing aggregate data about a large number of rounds, rather than about only one round, could lead to different results, because it would reduce path dependence. With respect to theory, we believe that more complex models, which combine conditional social preferences with self-confirming equilibrium, might improve our understanding of subjects' behavior. Our model was a first step in understanding how heterogeneous self-confirming beliefs interact with non-selfish preferences. Combining the framework of DFL with a model like the one introduced in Levine (1998) might be a difficult but worthwhile endeavor.

A Appendix: Proofs

Proof of Proposition 1

A simple dominance argument concerning the behavior of altruistic player 2 shows that in all THSCE where $\sigma_1(T_1) \neq 1$, it holds that $\sigma_2(T_2|\theta_2^2) = 0$, and in all THSCE where $\sigma_1(P_1P_3) > 0$, it holds that $\sigma_2(P_2P_4|\theta_2^2) = 1$. (Note that these conditions are also true in the setting with aggregate feedback). Moreover, in all THSCE where $\sigma_1(T_1) \neq 1$, it holds that $\sigma_2(P_2P_4|\theta_2^1) = 0$. To see this, assume on the contrary that there is a THSCE σ where $\sigma_1(T_1) \neq 1$ and $\sigma_2(P_2P_4|\theta_2^1) > 0$, so that condition 6 of Lemma 1 holds. If, for this THSCE, $\sigma_1(P_1P_3) = s_1(P_1P_3|\sigma_1) = 0$, then, for any conjectures $\widehat{\sigma_1}$ that satisfy condition 6, $s_1(P_1P_3|\widehat{\sigma_1}) = s_1(P_1P_3|\sigma_1) = 0$. Thus, $Eu_2(P_2P_4, \widehat{\sigma_1}) = 0.15 \cdot s_1(T_1|\sigma_1) + 0.6 \cdot [1 - s_1(T_1|\sigma_1)] < Eu_2(T_2, \widehat{\sigma_1}) = 0.15 \cdot s_1(T_1|\sigma_1) + 1.2 \cdot [1 - s_1(T_1|\sigma_1)]$, so P_2P_4 cannot be optimal. If $\sigma_1(P_1P_3) = s_1(P_1P_3|\sigma_1) > 0$, then for any conjectures that satisfy condition 6, $Eu_2(P_2P_4, \widehat{\sigma_1}) = 0.15 \cdot s_1(T_1|\sigma_1) + 0.6 \cdot s_1(P_1T_3|\sigma_1) + 2.4 \cdot s_1(P_1P_3|\sigma_1) < Eu_2(P_2T_4, \widehat{\sigma_1}) = 0.15 \cdot s_1(T_1|\sigma_1) + 0.6 \cdot s_1(P_1T_3|\sigma_1) + 4.8 \cdot s_1(P_1P_3|\sigma_1)$. We conclude that the condition $\sigma_2(P_2P_4|\theta_2^1) = 0$ must hold. (A very similar argument shows that this condition also holds in the presence of aggregate feedback. We will use the three conditions established in this paragraph in all three proofs.)

Let σ be a THSCE such that $\sigma_1(T_1) \neq 1$, which implies that $\pi_1(T_1|\sigma) \neq 1$. For any action that player 1 or selfish player 2 chooses with positive probability, the relevant condition of Lemma 1 must be satisfied. First, it is clear that $\pi_2(T_2|\sigma) \neq 1$, since $\sigma_2(T_2|\theta_2^2) = 0$. Now we claim that $\sigma_1(P_1P_3) \neq 0$. To show this, assume that, on the contrary, $\sigma_1(P_1P_3) = 0$.

Under this hypothesis, we argue that it must be the case that $\sigma_2(T_2|\theta_2^1) = 1$. If P_2T_4 was played with positive probability by selfish player 2, condition 5 of Lemma 1 would hold. There should then exist conjectures $\widehat{\sigma_1}$, such that they induce the correct distribution of actions about player 1, and such that the expected payoffs of action P_2T_4 are higher than the expected payoffs of action T_2, given the same conjectures. But it holds that, for any such conjectures, $Eu_2(P_2T_4, \widehat{\sigma_1}) = 0.15 \cdot s_1(T_1|\sigma_1) + 0.6 \cdot s_1(P_1T_3|\sigma_1) + 4.8 \cdot s_1(P_1P_3|\sigma_1) = 0.15 \cdot s_1(T_1|\sigma_1) + 0.6 \cdot [1 - s_1(T_1|\sigma_1)] < Eu_2(T_2, \widehat{\sigma_1}) = 0.15 \cdot s_1(T_1|\sigma_1) + 1.2 \cdot [1 - s_1(T_1|\sigma_1)]$.

Therefore, if the initial hypothesis (that $\sigma_1(P_1P_3) = 0$) were true, it would have to hold that $\sigma_2(T_2/\theta_2^1) = 1$, and, hence, $s_2(T_2|\sigma_2, p) = (1 - \xi)$. Since $\sigma_1(T_1) \neq 1$, and since we have assumed that $\sigma_1(P_1P_3) = 0$, P_1T_3 must be played with positive probability. Thus, condition 2 of Lemma 1 would have to hold. For any beliefs and conjectures that satisfy it, the expected payoffs of P_1T_3 are equal to $0.3 \cdot (1 - \xi) + 2.4 \cdot \xi$, and the expected payoffs of T_1 are 0.6. So, for P_1T_3 to be weakly preferred to T_1, it must be the case that $\xi \geq 1/7$, which is untrue.[46] Hence, the initial hypothesis must be wrong, so $\sigma_1(P_1P_3) > 0$. This implies that $\pi_3(T_3|\sigma) \neq 1$.

Since $\sigma_2(P_2P_4|\theta_2^2) = 1$, it holds that $\pi_4(T_4|\sigma) \neq 1$. Clearly, $\sigma \in \Sigma^*$, since $\sigma_1(P_1P_3) > 0$ and $s_2(T_2|\sigma_2, p) \neq 1$. Finally, it holds that $\pi_4(P_4|\sigma) = [\xi \cdot 1 + (1 - \xi) \cdot 0]/\pi_2(P_2|\sigma) \Rightarrow \xi = \pi_2(P_2|\sigma) \cdot \pi_4(P_4|\sigma)$. *QED*

Proof of Proposition 2

[46] Notice that any beliefs about the fractions of P_2T_4 and P_2P_4 can be used to rationalize the choice of action P_1T_3 as opposed to action P_1P_3.

Let σ be a THSCE such that $0 < \pi_i(T_i|\sigma) < 1$ for $i = 1, 2, 3, 4$. This profile must satisfy the relevant conditions of Lemma 1. Action P_2T_4 is played with positive probability by selfish 2, since $0 < \pi_4(T_4|\sigma)$, and $\sigma_2(P_2P_4|\theta_2^2) = 1$. So, condition 5 of the lemma must hold, and for any conjectures consistent with this condition, $Eu_2(P_2T_4, \widehat{\sigma_1}) = 0.15 \cdot s_1(T_1|\sigma_1) + 0.6 \cdot s_1(P_1T_3|\sigma_1) + 4.8 \cdot s_1(P_1P_3|\sigma_1)$, and $Eu_2(T_2, \widehat{\sigma_1}) = 0.15 \cdot s_1(T_1|\sigma_1) + 1.2 \cdot [s_1(P_1T_3|\sigma_1) + s_1(P_1P_3|\sigma_1)]$. Thus, since P_2T_4 is weakly optimal, $s_1(P_1T_3|\sigma_1) \leq 6s_1(P_1P_3|\sigma_1)$, and it follows that $\pi_3(P_3|\sigma) = s_1(P_1P_3|\sigma_1)/[s_1(P_1T_3|\sigma_1) + s_1(P_1P_3|\sigma_1)] \geq 1/7$.

Similarly, since P_1P_3 is played with positive probability, condition 3 of the lemma holds. For any beliefs and conjectures that satisfy condition 3, $Eu_1(P_1P_3, \widehat{\sigma_2}, \widetilde{\mu_1}) = 0.3 \cdot s_2(T_2|\sigma_2, p) + 1.2 \cdot s_2(P_2T_4|\sigma_2, p) + 9.6 \cdot s_2(P_2P_4|\sigma_2, p)$, and $Eu_1(P_1T_3, \widehat{\sigma_2}, \widetilde{\mu_1}) = 0.3 \cdot s_2(T_2|\sigma_2, p) + 2.4 \cdot [s_2(P_2T_4|\sigma_2, p) + s_2(P_2P_4|\sigma_2, p)]$. Since P_1P_3 is weakly optimal, $s_2(P_2T_4|\sigma_2, p) \leq 6s_2(P_2P_4|\sigma_2, p)$, so $\pi_4(P_4|\sigma) = s_2(P_2P_4|\sigma_2, p)/[s_2(P_2T_4|\sigma_2, p) + s_2(P_2P_4|\sigma_2, p)] \geq 1/7$.[47]

Moreover, $\sigma_1(P_1T_3) > 0$, so condition 2 of the lemma holds. For any beliefs and conjectures that satisfy condition 2, $Eu_1(P_1T_3, \widehat{\sigma_2}, \widetilde{\mu_1}) = 0.3 \cdot s_2(T_2|\sigma_2, p) + 2.4 \cdot [1 - s_2(T_2|\sigma_2, p)]$ and $Eu_1(T_1, \widehat{\sigma_2}, \widetilde{\mu_1}) = 0.6$, so it must be that for P_1T_3 to be optimal, $s_2(T_2|\sigma_2, p) \leq 6/7$. But notice that $\pi_2(P_2|\sigma) = 1 - s_2(T_2|\sigma_2, p)$, so $\pi_2(P_2|\sigma) \geq 1/7$. Note that we need not also compare the expected payoffs of action P_1T_3 with the expected payoffs of action P_1P_3, since the choice of the former action can be rationalized by any beliefs about the probabilities of actions P_2T_4 and P_2P_4.

Similarly, note that the choice of actions T_1 and T_2 is easy to rationalize by arbitrary beliefs about opponents' play, and hence does not impose additional constraints. Since $\xi = \pi_2(P_2|\sigma) \cdot \pi_4(P_4|\sigma)$, the conditions above imply that $\pi_2(P_2|\sigma) \leq 7\xi$ and $\pi_4(P_4|\sigma) \leq 7\xi$. *QED*

Proof of Proposition 3

Consider a strategy profile σ such that $\sigma_1(T_1) \neq 1$. We shall examine which conditions have to hold, in order for σ to be a THSCE with aggregate information. Since $\sigma_2(T_2|\theta_2^2) = 0$, it holds that $s_2(T_2|\sigma_2, p) \neq 1$, so the aggregate signal function of population 2 ensures that the conjectures of all individual 2's induce a distribution of 1's actions which is the same as the actual distribution. So the subjective expected payoffs from each action, T_2, P_2T_4, and P_2P_4, coincide with the true expected payoffs.

If it was true that $\sigma_1(P_1P_3) = 0$, then the best response of selfish individual 2's would be T_2 (since $1.2 > 0.6$). So, it would have to hold that $s_2(T_2|\sigma_2, p) = (1 - \xi)$. Notice that P_1T_3 should be necessarily played with positive probability, since $\sigma_1(T_1) \neq 1$. The signal function of population 1 would then imply that the true fraction $s_2(T_2|\sigma_2, p)$ is expected. Then, for P_1T_3 to be optimal, it would have to hold that $\xi \geq 1/7$ (see proof of Proposition 1), which is untrue. Therefore, we conclude that $\sigma_1(P_1P_3) \neq 0$.

Now assume that $\sigma_1(P_1P_3) = 1$, (so that all individual 1's have correct beliefs about the distribution of actions of population 2). The best reply of selfish 2's is then P_2T_4 (since $4.8 > 2.4 > 1.2$), which implies that $s_2(P_2P_4|\sigma_2, p) = \xi, s_2(P_2T_4|\sigma_2, p) = (1 - \xi)$. So the

[47] Action P_1P_3 is weakly optimal, since in the next paragraph we also show that $Eu_1(T_1, \widehat{\sigma}_2, \widetilde{\mu}_1) < Eu_1(P_1T_3, \widehat{\sigma}_2, \widetilde{\mu}_1)$.

expected payoffs of P_1P_3 are equal to $9.6 \cdot \xi + 1.2 \cdot (1-\xi)$ and the expected payoffs of P_1T_3 are equal to 2.4. Hence, P_1P_3 cannot be optimal, since $\xi < 1/7$. We conclude that our initial assumption cannot be true, so $\sigma_1(P_1P_3) \neq 1$.

Similarly, if it was the case that $\sigma_1(P_1T_3) = 0$, then, since P_1P_3 cannot be played with probability one, both T_1 and P_1P_3 must be played with positive probability. Therefore, individual 1's would have to be indifferent between T_1 and P_1P_3. But since the best response of all selfish 2's would then be P_2T_4 (since $4.8 > 2.4 > 1.2$), individual 1's should strictly prefer action P_1P_3 to action T_1, because $1.2 \cdot (1-\xi) + 9.6 \cdot \xi > 0.6$.

Consequently, both P_1T_3 and P_1P_3 must be played with positive probability, so the expected payoffs of P_1T_3 and P_1P_3 must be the same. Hence, $2.4[s_2(P_2T_4|\sigma_2, p)+s_2(P_2P_4|\sigma_2, p)] = 1.2 \cdot s_2(P_2T_4|\sigma_2, p) + 9.6 \cdot s_2(P_2P_4|\sigma_2, p)$ so that $s_2(P_2P_4|\sigma_2, p) = s_2(P_2T_4|\sigma_2, p)/6$. Therefore, $\pi_4(P_4|\sigma) = s_2(P_2P_4|\sigma_2, p)/[s_2(P_2P_4|\sigma_2, p) + s_2(P_2T_4|\sigma_2, p)] = 1/7$.

Remember that $\sigma_2(P_2P_4|\theta_2^1) = 0$. Since $\pi_4(P_4|\sigma) = [\xi \cdot 1 + (1-\xi) \cdot 0]/[1-(1-\xi) \cdot \sigma_2(T_2|\theta_2^1)] = 1/7 > \xi$, it must be the case that $\sigma_2(T_2|\theta_2^1) > 0$. Since $\pi_4(P_4|\sigma) \neq 1$, it holds that $\sigma_2(P_2T_4|\theta_2^1) > 0$. This implies that the expected payoffs, for selfish individual 2's, of action P_2T_4 are equal to the expected payoffs of action T_2. Therefore, it holds that $1.2 \cdot [s_1(P_1P_3|\sigma_1) + s_1(P_1T_3|\sigma_1)] = 0.6 \cdot s_1(P_1T_3|\sigma_1) + 4.8 \cdot s_1(P_1P_3|\sigma_1) \Longrightarrow 6s_1(P_1P_3|\sigma_1) = s_1(P_1T_3|\sigma_1)$. Accordingly, $\pi_3(P_3|\sigma) = s_1(P_1P_3|\sigma_1)/[s_1(P_1T_3|\sigma_1) + s_1(P_1P_3|\sigma_1)] = 1/7$.

Moreover, since $1/7 = \pi_4(P_4|\sigma) = \xi/\pi_2(P_2|\sigma)$, it follows that $\pi_2(P_2|\sigma) = 7\xi$. Finally, since the expected payoffs of action T_1 are lower than the expected payoffs of actions P_1P_3 and P_1T_3 (which are equal to $14.7\xi + 0.3$), we conclude that $\pi_1(T_1|\sigma) = 0$. *QED*

References

Olivier Armantier. Does observation influence learning? *Games and Economic Behavior*, 46 (2):221–239, 2004.

Pierpaolo Battigalli. Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali. Master's thesis, Bocconi University, 1987.

Pierpaolo Battigalli and Martin Dufwenberg. Dynamic psychological games. *Journal of Economic Theory*, 144(1):1–35, 2009.

Joyce Berg, John Dickhaut, and Kevin McCabe. Trust, reciprocity, and social history. *Games and Economic Behavior*, 10(1):122–142, 1995.

Robert Cialdini and Noah Goldstein. Social influence: Compliance and conformity. *Annual Review of Psychology*, 55:591–621, 2004.

Ralph D' Agostino, Warren Chase, and Albert Belanger. The appropriateness of some common procedures for testing the equality of two independent binomial populations. *The American Statistician*, 42(3):198–202, 1988.

Eddie Dekel, Drew Fudenberg, and David K. Levine. Learning to play bayesian games. *Games and Economic Behavior*, 46(2):282–303, 2004.

Martin Dufwenberg and Uri Gneezy. Information disclosure in auctions: an experiment. *Journal of Economic Behavior and Organization*, 48(4):431–444, 2002.

Mark Fey, Richard D McKelvey, and Thomas R Palfrey. An experimental study of constant-sum centipede games. *International Journal of Game Theory*, 25(3):269–87, 1996.

Bruno Frey and Stephan Meier. Social comparisons and pro-social behavior: Testing 'conditional cooperation' in a field experiment. *American Economic Review*, 94(5):1717–1722, 2004.

Drew Fudenberg and David K Levine. Self-confirming equilibrium. *Econometrica*, 61(3): 523–45, 1993.

Drew Fudenberg and David K Levine. Measuring players' losses in experimental games. *The Quarterly Journal of Economics*, 112(2):507–36, 1997.

Francis Fukuyama. Social capital, civil society and development. *Third world quarterly*, 22 (1):7–20, 2001.

Francis Fukuyama. Social capital and development: the coming agenda. *SAIS Review*, 22 (1):22–37, 2002.

Glenn Harrison and Kevin McCabe. Expectations and fairness in a simple bargaining experiment. *International Journal of Game Theory*, 25(3):303–327, 1996.

Matthew Jackson and Ehud Kalai. Social learning in reccuring games. *Games and Economic Behavior*, 21(1):102–134, 1997.

Ehud Kalai and Ehud Lehrer. Rational learning leads to nash equilibrium. *Econometrica*, 61(5):1019–45, 1993.

Toshiji Kawagoe and Hirokazu Takizawa. Level-k analysis of experimental centipede games. Ssnr paper, 2008. URL `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1289514`.

David K. Levine. Modeling altruism and spitefulness in experiment. *Review of Economic Dynamics*, 1(3):593–622, 1998.

Richard McKelvey and Thomas Palfrey. Quantal response equilibria for extensive form games. *Experimental Economics*, 1(1):9–41, 1998.

Richard McKelvey, Andrew McLennan, and Theodore Turocy. Gambit:software tools for game theory, 2007. Version 0.2007.01.30, 2007. http://www.gambit-project.org.

Richard D McKelvey and Thomas R Palfrey. An experimental study of the centipede game. *Econometrica*, 60(4):803–36, 1992.

Ryan Murphy, Amnon Rapoport, and James Parco. The breakdown of cooperation in iterative real-time trust dilemmas. *Experimental Economics*, 9(2):147–166, 2006.

Rosemarie Nagel and Fang Fang Tang. Experimental results on the centipede game in normal form: an investigation on learning. *Journal of Mathematical Psychology*, 42:356–384, 1998.

Andreas Ortmann, John Fitzgerald, and Carl Boeing. Trust, reciprocity, and social history: A re-examination. *Experimental Economics*, 3:81–100, 2000.

Ignacio Palacios-Huerta and Oscar Volij. Field centipedes. *American Economic Review*, 99 (4):1619–1635, 2009.

Kristopher Preacher and Nancy Briggs. Calculation for fisher's exact test: An interactive calculation tool for fisher's exact probability test for 2 x 2 tables, 2001. Available from http://quantpsy.org.

Amnon Rapoport, William E. Stein, James E. Parco, and Thomas E. Nicholas. Equilibrium play and adaptive learning in a three-person centipede game. *Games and Economic Behavior*, 43(2):239–265, 2003.

Hardeo Sahai and Adwer Khurshid. On analysis of epidemiological data involving a 2x2 contingency table: An overview of fisher's exact test and yates' correction for continuity. *Journal of Biopharmaceutical Statistics*, 5(1):43–70, 1995.

Klaus G. Zauner. A payoff uncertainty explanation of results in experimental centipede games. *Games and Economic Behavior*, 26(1):157–185, 1999.