



OPEN

Inferring strategies from observations in long iterated Prisoner's dilemma experiments

Eladio Montero-Porras^{1,5✉}, Jelena Grujić^{1,5}, Elias Fernández Domingos^{1,2} & Tom Lenaerts^{1,2,3,4}

While many theoretical studies have revealed the strategies that could lead to and maintain cooperation in the Iterated Prisoner's dilemma, less is known about what human participants actually do in this game and how strategies change when being confronted with anonymous partners in each round. Previous attempts used short experiments, made different assumptions of possible strategies, and led to very different conclusions. We present here two long treatments that differ in the partner matching strategy used, i.e. fixed or shuffled partners. Here we use unsupervised methods to cluster the players based on their actions and then Hidden Markov Model to infer what the memory-one strategies are in each cluster. Analysis of the inferred strategies reveals that fixed partner interaction leads to behavioral self-organization. Shuffled partners generate subgroups of memory-one strategies that remain entangled, apparently blocking the self-selection process that leads to fully cooperating participants in the fixed partner treatment. Analyzing the latter in more detail shows that ALLC, ALLD, TFT- and WSLS-like behavior can be observed. This study also reveals that long treatments are needed as experiments with less than 25 rounds capture mostly the learning phase participants go through in these kinds of experiments.

The emergence of cooperation in human populations has been an object of study in many different disciplines ranging from social sciences^{1–3}, physics and complex systems^{4,5} to biology^{6,7} and computer science⁸, among many others. While cooperation is widespread in human societies, its origins remain unclear, although many mechanisms have been proposed to explain its evolutionary advantage². In fact, many social situations pose a dilemma, in which, while being cooperative towards society is beneficial for the population as a whole, free-riding on the efforts of others may generate substantial individual gains^{9,10}. While many studies have explored which individual behaviors can promote cooperation in these social dilemmas^{8,11,12} or aimed to find how much human decision-making aligns with those cooperative strategies^{1,13}, not so much is known about which strategies effectively can be extracted from experimental game-theoretical data, which often do not end in full cooperation. Identifying the actual decision-making schemes is essential, especially if one wants to understand why cooperation is not achieving anticipated levels, to grasp the incentives needed to promote beneficial outcomes or how artificial systems may need to be designed so that they align with human pro-social behavior.

Using a data science approach we aim to provide an answer to such questions here, focusing on a new cohort of game theoretical experiments within the well-known framework of the pairwise Iterated Prisoner's dilemma (IPD)¹⁴. In the IPD, when both players cooperate (C) in a round, they both get a reward (represented by R), while if one of them defects (D), and the other cooperates they get a payoff T and S respectively. If both defect, they both get a payoff P . The dilemma emerges when $T > R > P > S$, with $2R > T + S$. The payoff is accumulated by the interacting players at each round until the game ends. In this iterated version of the one-shot PD there are a plethora of possible equilibrium outcomes¹⁵, including cooperative ones¹⁶.

The literature has reported on many strategies that can induce cooperation in the IPD with fixed partners. These results were produced mainly through models and simulations. In the famous first Axelrod tournament¹⁷, for instance, Tit-for-Tat (TFT) emerged as the most successful strategy. TFT starts by signaling the intention to cooperate and afterwards mimics the previous action of the opponent, and it has been argued to be a good representation of reciprocal behavior in human societies. Reciprocity has been thoroughly researched as one of the most important mechanisms to favour cooperation^{1,2,18}. More generally, conditionally cooperative strategies

¹Artificial Intelligence Laboratory, Vrije Universiteit Brussel, Brussels 1050, Belgium. ²Machine Learning Group, Université Libre de Bruxelles, Brussels 1050, Belgium. ³Center for Human-Compatible AI, UC Berkeley, Berkeley 94702, USA. ⁴FARI Institute, Université Libre de Bruxelles-Vrije Universiteit Brussel, 1050 Brussels, Belgium. ⁵These authors contributed equally: Eladio Montero-Porras and Jelena Grujić. ✉email: eladio.montero@vub.be

appear to generate the right set of opportunities for cooperative behavior to spread in large populations^{19–21}, highlighting the importance of context and past experiences in the effectiveness of cooperative strategies. Reciprocity thus justifies, from an evolutionary perspective, the existence of altruistic and pro-social behaviors²².

However successful TFT was to explain aspects such as direct reciprocity and conditional cooperation, this strategy is known to dissolve into mutual defection in the presence of execution errors, i.e. when participants fail to keep the implicit cooperative agreement at a given round²³, especially in heterogeneous populations²⁴. Attempts to repair this flaw lead to the introduction of Generous-Tit-for-Tat (GTFT)²⁵ which cooperates if the opponent cooperates, but if the opponent defects it sometimes “forgives” and continues to cooperate. This strategy is also a reactive strategy, but unlike TFT, it is stochastic, because the player’s next action is now given with some probability. Yet, GTFT has its own problems leading to the introduction of the Win-Stay-Lose-Shift (WSLS)²⁶ strategy that repeats the action from the previous round if the player is happy with the obtained payoff (T and R), otherwise changes to the opposite action (when the payoff is P or S). More strategies have been proposed and analyzed since then (see Martinez-Vaquero et al²⁷ for a comprehensive study).

From an experimental angle, only a few works have examined how participant actions in the IPD may be translated into relevant strategies used by humans, while also studying what factors may affect the observed behavior. Some experiments with animals showed that guppies⁷, sticklebacks²⁸ and tree swallows²⁹, exhibit a TFT-like behavior. In human behavioral economic experiments, both the theoretical GTFT and WSLS appear to align with the decision-process of the subjects³⁰. Dal Bó et al³¹ showed in another IPD experiment that, when given the choice among theoretical strategies, subjects preferred strategies like Always Defect (AllD), TFT, and Grim Trigger (starts cooperating, cooperates as long as both players have cooperated in the last round, and defects otherwise) and did not use WSLS, which was shown to have more “desirable” properties, such as not defecting forever after a deviation.

While matching theoretical strategies to experiments or asking people to select their preferred strategy provides insight into how they relate to human preferences, it does not immediately reveal how humans actually decide to act in the IPD while playing freely, since their strategies can change with time and respond to different factors, such as their opponents’ actions or simply by learning the game and reach their own equilibrium. Alternatively, their strategy may be directly inferred from the data itself, using algorithmic models³². This approach was taken to determine the strategies in an Ultimatum game experiment, where binary decision trees were used to model the strategies³³ or symbolic regression³⁴. A similar approach was taken to infer the behavior of participants in Trust games^{35,36} or in a market simulation, where strategies were inferred using Bayesian inference^{37,38}. Yet so far, no inference of strategic models has been performed on IPD, a caveat this work is overcoming.

We present thus, on the one hand, the results of a behavioral economics experiment that investigates how subjects act in the IPD and, on the other hand, the strategies that can be extracted from the data. The experiment consists of two treatments: i) one where the subjects play the IPD with a fixed partner over a large, unspecified number of rounds, which will be called the *fixed partners* or *FP* treatment; and ii) a second long IPD treatment where the opponent changes each round, which we call *shuffled partners* or *SP* treatment. The objective of collecting the data over a large number of rounds was to understand what the effect is of a long experiment on the level of cooperation in the IPD and to study how the inferred behavior differs over time and between a setting where partnerships can be established or one where one is repeatedly confronted by strangers^{39–41}.

It is not clear from the literature how many rounds of the IPD are needed to observe a stabilization in the human decision-making process, meaning that the learning phase has passed and people are acting according to a clearly defined strategy. Some works have studied the strategies subjects play for ten rounds⁴², another one that studied the evolution of cooperation in the IPD lasted on average 1.96 and 4.42 rounds, for their two treatments⁴³, another work studied strategies in the IPD with noise lasted on average 8 rounds, where they noticed a considerable strategic diversity, suggesting the subjects did not learn the game completely⁴⁴, others, a range of 10 to 35 rounds⁴⁵, 15 rounds⁴⁶ and 100 rounds⁴⁷, to cite a few. Yet, the latter focuses on an interaction with a fixed theoretically defined agent. To study the evolution in the decision-making process, we need to know how long the subjects need to learn the game among themselves in different points in time. Given this insight, one can then use the algorithmic modeling approach to infer the actual behaviors, comparing them over time and over different treatment settings, as is the case here.

As experience determines the behavior of each participant, which we will call context, we will use unsupervised clustering techniques to identify the contexts that are found in both IPD treatment data and then determine which strategic algorithmic models may be inferred from the data in each of the contexts (with the context being defined by their own actions and their opponents’ in the previous round). We focus in this work on memory-one strategies since the experiments only informed participants of what happened in a previous round. Moreover, it has been shown that for indefinitely repeated games, players’ payoffs with memory-one strategies are the same regardless if their co-players use longer-memory strategies⁴⁸, making it difficult to discern the significance of the longer-memory ones.

Our hypothesis is that different behaviors will be inferred from the differing context experiences. Iterating the game for a sufficient time, as specified earlier, will thus allow to clearly discern their short and long-term behavior. We designed the experiment to last 100 rounds, without explicitly informing this hard limit to the participants, which is much longer than previous experiments^{43–45} to cite a few, and investigated how many rounds are actually needed before the learning process ends and the participant behavior appears to become consistent. In other long experiments, this stabilization seemed to appear after 10–20 rounds⁴⁹, which will be examined in more in a detailed analysis here.

The participant behavior is modeled using Hidden Markov Models (HMM)⁵⁰, and its parameters are trained using `hmmlearn`⁵¹ on the treatment data separated in behavioral clusters, while simultaneously trying to find the preferred minimal HMM structure (number of states and transition structure) to achieve this. The resulting HMM models are both simple and transparent, containing enough modeling power to represent subjects’

	C	D
C	3,3	0,4
D	4,0	1,1

Table 1. Payoff matrix for both IPD treatments.

strategies in the IPD while also being generative. The latter is of interest as the inferred strategies could be directly be used within the context of theoretical simulations assessing their performance, which is left for future work.

Methods

Experimental data collection. As explained in the Introduction, data was collected for two treatments wherein participants played long IPD in two different pairwise configurations, i.e. fixed partners (FP) and shuffled partners (SP). The data from these experiments were collected in Brussels, Belgium, at the Brussels Experimental Economics Laboratory (BEEL), part of the Vrije Universiteit Brussel (VUB). All experiments followed the relevant guidelines and regulations of data protection and experiments with human participants and were approved by the Ethical Commission for Human Sciences at the VUB (ECHW2015_3). Moreover, all participants gave consent to the experiment by signing a consent form after the instructions of the experiment were read and all questions the participants had were answered and addressed.

Table 1 shows the payoff matrix containing the per round rewards used in the two treatments. For both treatments, participants could observe the actions of their partner in the previous round, even when this partner changed from the previous to the current round. More information about the experimental sessions and details about the data collected can be found in the Supplementary Information.

For each participant in both treatments, we collected their choices, as is visualized by panel A in Fig. 1. The combination of two actions, i.e. *CC*, *CD*, *DC* or *DD* (action format: player-opponent, e.g. *CD* means the focal player cooperated and their opponent defected in the previous round), provides a context for the next round, as participants will get this information when making their decision. Each context can now be combined with the action of a participant after observing that context and this combination, as shown in panel B of Fig. 1, can now be translated into one of eight values. The entire sequence of actions of a participant and the associated contexts can thus be transformed into a new sequence of numbers that represent the *conditional actions* of each participant. This new sequence will be used to train an HMM, representing in its emission probabilities in every state the conditional response, i.e. cooperate or defect, based on the context they experienced. We also collect for each context, the probability of observing a conditionally cooperative action, as visualized in Figs. 4 and 6 in the Results section.

Clustering contexts and context-dependent behaviors. People act according to their preferences, which includes how they think others should behave. In order to understand the strategies that are being used one needs to explicitly consider the context wherein they happen. Thus, to find the strategies in the FP and SP treatments, we first separate participants according to their experiences. Once this is complete one can question whether each participant displays the same response for the contexts they experience. This allows one to correctly grasp how humans act in the IPD. As mentioned in the introduction, we focus in this work on memory-one strategies. There are two reasons for this focus: 1) as Press et al. mentioned, in indefinitely repeated games memory-one strategies have the same payoff against those with a longer memory⁴⁸, making it problematic to discern them. This argument holds for our experiments as the participants did not know how many rounds the experiment would take. 2) The experimental interface (see Supplementary Information) only reminded participants of their and their opponent's actions in the last round. It may be that participants remember actions from two or more rounds in the past, yet they were not triggered to do so. Our methodology nevertheless allows for the identification of strategies with longer memories, but we have left that for future work as new experiments and more extensive analysis would be required to estimate the importance of each longer-memory strategy against the memory-one's that are studied here.

To identify the context groups and the behaviors within those groups, a cluster analysis was performed here: we first clustered the subjects based on the number of times of (*CC*), (*CD*), (*DC*) and (*DD*) happened, providing a *contextual clustering*. Subsequently, another clustering was performed on the number of cooperative actions in the sequence for each of the previous action combinations, i.e. (*CC*)C, (*CD*)C, (*DC*)C and (*DD*)C, providing thus a *behavioral clustering per context*. These variables are used to generate the t-distributed stochastic neighbor embedding (t-SNE)⁵² plot in Fig. 3 and Supplementary Fig. S3.

Different clustering approaches were considered (e.g. K-means, Hierarchical Clustering, and Network Modularity analysis, as can be observed in Supplementary Information, see section 1.3 for more information on the Modularity Network clustering) but in the end, K-means is sufficient as the other approaches generated similar results (see for Supplementary Fig. S1, S3 top row, for a comparison with Hierarchical Clustering and Fig. S4 for Modularity Network Clustering). The implementations were done in Scikit learn⁵³. To determine the optimal number of clusters, we used the “elbow” method by Santopaa et al.⁵⁴ which plots a curve with the sum of squared distances of samples to their closest cluster center and chooses the minimum best number of clusters that minimize the inertia. See Supplementary Figs. S1 and S2 for the results produced but the “elbow” method.

To evaluate the quality of the results generated by clustering algorithms, we use the silhouette measure⁵⁵, which measures the mean intra-cluster Euclidean distance and the mean nearest-cluster Euclidean distance for

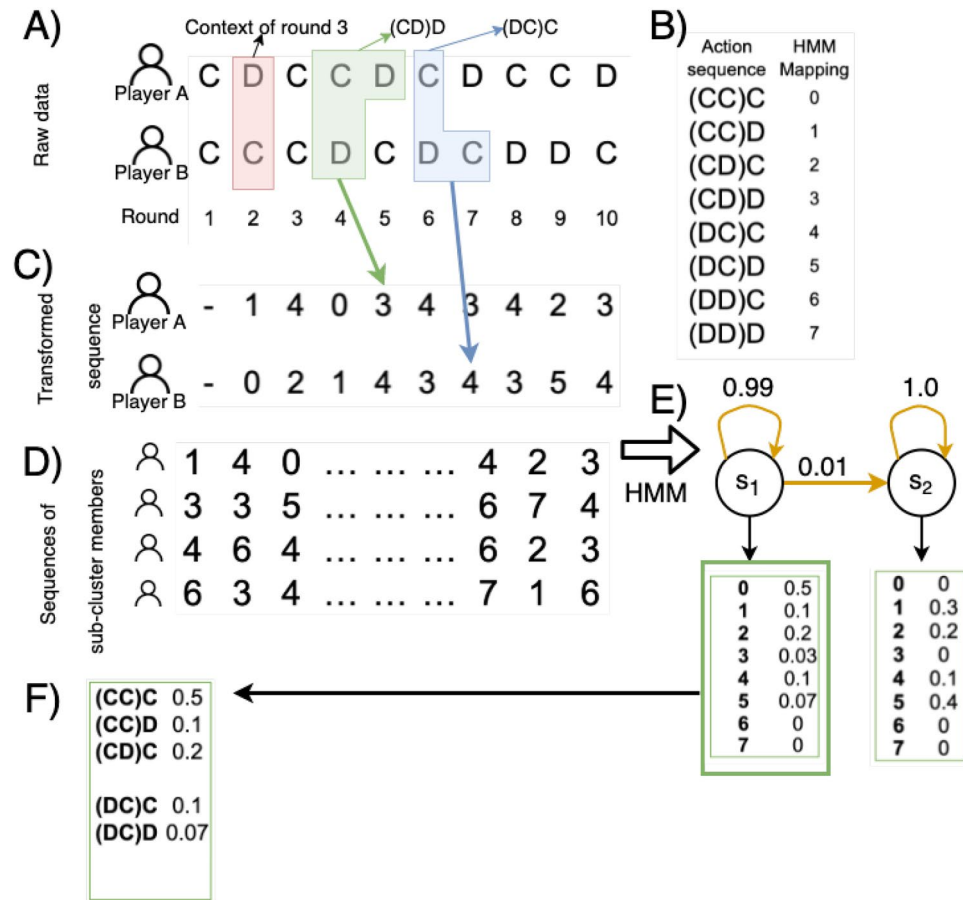


Figure 1. From treatment data to HMM. (A) At the top, the first 10 rounds between a player and her opponent, in the red box, the actions of round 2, which represent the context for round 3. (B) The action sequence mapping, with the green box in panel A representing the action sequence (CD)D for player A and the blue box (DC)C also in panel A for player B in rounds 5 and 7 respectively. (C) The transformed sequence for players A and B into integers so the HMM with the mapping table in panel B. (D) The resulting sequences are grouped given their sub-cluster and then processed by the HMM in panel (E). In this example, the resulting HMM is a model with two states s_1 and s_2 with transition probabilities given by the yellow arrows with emission probabilities in the boxes below. (F) HMM visualization. Note that all the emission probabilities ≤ 0.05 and all transition probabilities ≤ 0.01 were not taken into account for readability. The resulting HMM emission table is colored green from the hidden state s_1 .

each observation, with 1 is being the best score and values near zero indicate that some clusters might be overlapping, while negative values indicate observations assigned to a wrong cluster.

Inferring Hidden Markov Models. Given the context clusters and the context-dependent behavioral sub-clusters, a HMM, using the conditional action sequences (see Fig. 1), is produced for each sub-cluster. As shown in the figure, a HMM is composed of a number of states (s_1 and s_2 in the figure) which are connected by transitions (yellow arrows). The HMM is a probabilistic Markov Model in which each observation of a sequence is produced by a hidden (non-observable) state⁵⁰. We used a multinomial model for HMM with a sequential structure, this means that the model has states connected from left to right, where a transition to another state can only be made in that direction. Returns to a previous hidden state are thus not possible. As also shown in the figure, the transformed sequence of conditional actions is used to train the HMM. The procedure to determine the optimal number of hidden states is specified in Algorithm 1 in the Supplementary Information.

To train the HMM, the `hmmlearn` library⁵¹ was used. To visualize the resulting models the `GraphViz` package for Python⁵⁶ was used. We expect the participant choices to be stochastic since subjects were not instructed to use any strategy in particular, but the action they considered was best according to their expectations. For this reason, we expect the inferred strategies to be noisy. For visualization, we mapped the number sequences back to human-readable triplets, and all the emission probabilities in the HMM less than 0.05 were discarded, as shown in Fig. 1*.

Evolution of the strategies. To analyze how the players changed and adapted their strategy over all the 100 rounds, the same procedure of the two-fold clustering and HMM modeling was performed on four differ-

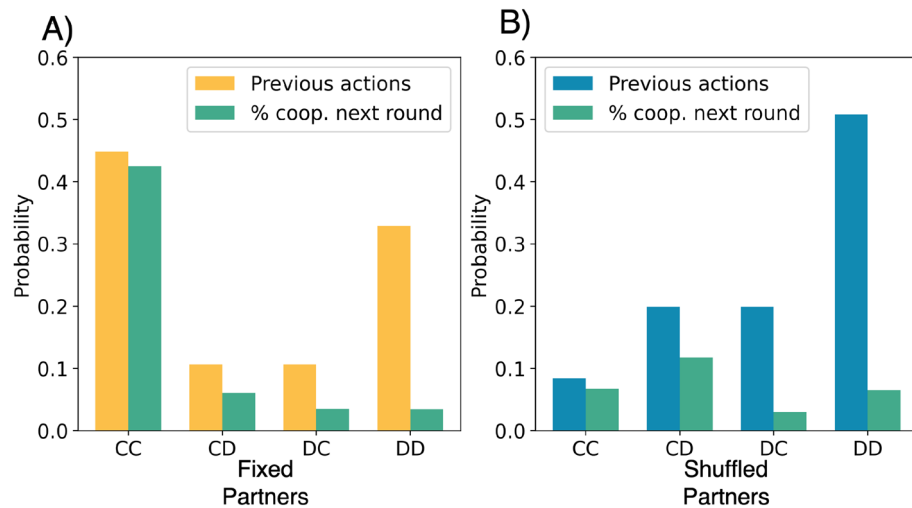


Figure 2. Probability of each context and the cooperation rate per treatment. In the x-axis, the context of each decision is shown, i.e. the actions of the previous round. The green bar in each chart shows the fraction of subsequent cooperative actions when experiencing the particular context.

ent round windows, i.e. from round 1-25, 26-50, 51-75, and finally rounds 76-100). The objective is to assess whether the strategy stabilizes over time and to examine how the treatment, i.e. fixed or shuffled partners, affect the results.

Ethical approval. Ethical approval by reference number ECHW2015_3 was obtained from the Ethical Commission for Human Sciences at the Vrije Universiteit Brussel to perform this experiment. All experiments were performed in accordance with the European Union GDPR guidelines and regulations, and the study was conducted in accordance with the Declaration of Helsinki. All the participants in the study had to give their informed consent prior to the participation. All the data of the experiment has been anonymized and cannot be linked to any participant.

Results

Specific contextual subgroups with associated behavioral responses emerge in each treatment. Before identifying the different clusters in the data provided by both treatments (see “Methods” section and Supplementary Information for details), we first examine the distribution of contexts experienced by participants during the experiment. Figure 2 shows this distribution while also revealing the fraction of times each context led to a cooperative response by a participant. In the case of the FP treatment, mutual cooperation (CC) or mutual defection (DD) is observed most often and was matched with the expected response (i.e. either $(CC) \rightarrow C$ or $(DD) \rightarrow D$, which we write as $(CC)C$ and $(DD)D$ respectively). In the SP treatment, mutual defection DD and the anticipated response D was most prevalent. Although a similar cooperation probability was shared by the two treatments (see Supplementary Table S3) in the case of DC and CD, there was a higher probability of cooperation in FP than in SP in the DC context, i.e. responding positively to an act of cooperation of the co-player. Overall, the figure shows that by shuffling partners cooperation is reduced, as anticipated by prior experiments³⁹. Also, it already hints at different clusters that could be inferred from both data sets.

Clustering the treatments on the contextual information that each participant experiences (the frequency of contexts CC, CD, DC and DD) with K-means (see “Methods” section) singled out 3 groups in both FP and also 3 in SP. As reported in “Methods” section and visualized in Supplementary Fig. S1, different algorithms (network modularity and hierarchical clustering) were tested, revealing that a similar number of clusters was obtained in each case. Cluster quality is assessed using silhouette scores, which are 0.5445 and 0.4201 respectively for the FP and SP treatment. These results indicated that a relatively good separation is found between the different contextual clusters.

Panel A in Fig. 3 shows the composition of each cluster in terms of its context. In the FP treatment (A, top row), three different groups are identified: Cluster A containing the players who experience mutual cooperation CC most often, cluster B where they experience mutual defection DD most, and then cluster C where experiences are mixed. These results were anticipated, given the differences shown in Fig. 2. In SP (bottom row Fig. 3A) mutual defection is favored over mutual cooperation in two of the clusters, i.e. clusters D and E. Nonetheless, there are sufficient differences between them which are captured by our clustering approach.

The tSNE plot (see “Methods” section) in panel B of Fig. 3 allows for the exploration of these differences and compares the clusters identified in both treatments. This plot reveals that clusters A and C, which are more cooperative in experienced contexts and responses, clearly differentiate themselves from the rest, and one can find them at completely the other side of experiences and responses to those belonging to cluster B of the FP treatment. There seems nonetheless to be an overlap between some members cluster C and those in F in the

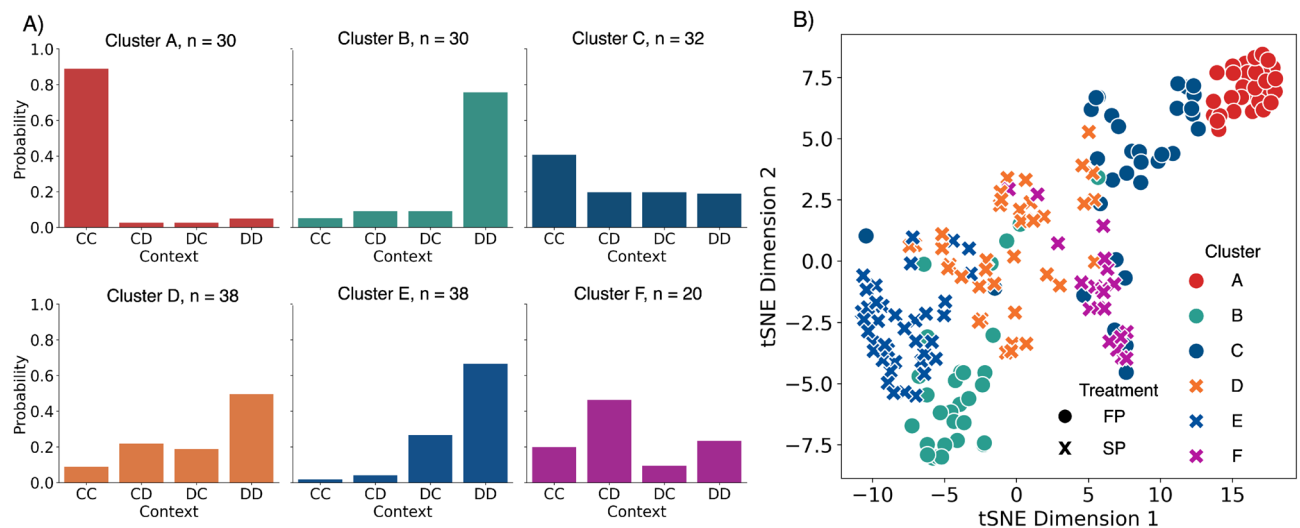


Figure 3. Probability of cooperation given each context per cluster (using K-Means). (A) The top row shows the context composition of each cluster in the FP treatment, the bottom row does the same for each SP cluster. (B) tSNE visualization of the different clusters A–F.

Treatment	Context	No. behavioral	Silhouette
	Cluster	Sub-clusters	Score
FP	A ($n = 30$)	2	0.6868
FP	B ($n = 30$)	3	0.3848
FP	C ($n = 32$)	3	0.4627
SP	D ($n = 38$)	4	0.2993
SP	E ($n = 38$)	4	0.4417
SP	F ($n = 20$)	3	0.4176

Table 2. Summary of the clustering results using K-Means. The contextual clustering results in 3 clusters of similar sizes. The behavioral clustering per context yielded 8 behavioral sub-clusters in FP and 11 behavioral sub-clusters in SP. The scikit-learn⁵³ silhouette analysis was used to determine the quality of the clusters generated with the “elbow” method (see “Methods” section).

SP treatment, yet most C members form a group on themselves. The experiences and behaviors of members of the E cluster, on the other hand, are close to those in B, which one can also observe when comparing the two bar charts. Finally, cluster D consists of participants that have experiences and responses that lie in between all others, separating the more defective and cooperative spectrum.

Given these observations, in the following sections, we show how we can infer the behaviors adopted by participants in each contextual cluster. It is important to note that not considering explicitly the experiences of the participants and focusing directly on the conditional behavior may result in grouping together participants that have encountered a different distribution of contexts. It is therefore essential to first partition the participants in function of their contexts distribution and then determine how they differ in their choices given these experiences.

Fixed partners promotes behavioral self-selection to cooperative or defecting behaviors. Table 2 reports the number of behavioral sub-clusters for each contextual cluster (see “Methods” section): in total, 8 subgroups were inferred from the raw data of the FP treatment, i.e., 2 in cluster A, 3 in cluster B, and 3 in cluster C. Supplementary Fig. S3 shows a tSNE visualization of how the treatment data is divided into clusters and sub-clusters for FP. As before, the best settings for K in K-means clustering were produced using the “elbow” method⁵⁴ (see “Methods” section and Supplementary Fig. S2).

The behaviors in each contextual cluster are captured each time by two plots: i) a first plot that shows the likelihood of cooperation when experiencing a given context in a contextual cluster, and ii) a second plot that provides the inferred HMM in the behavioral sub-cluster (see “Methods” section).

For cluster A, one can observe the difference in response when it comes to a *DD* situation, when comparing both sub-clusters (results in red): Individuals in cluster A.0 are more likely to cooperate again, than those in A.1 (see Fig. 4). Moreover, the resulting HMM for cluster A.0 shows how an initial strategy of reciprocating defection ((*CD*)*D*), leads to mutual cooperation ((*CC*)*C*). In addition, an increase in forgiving defective behavior ((*CD*)*C*) can be observed in Cluster A.1 (see HMM for A.1 in Fig. 5), which is present less often in cluster A.0.

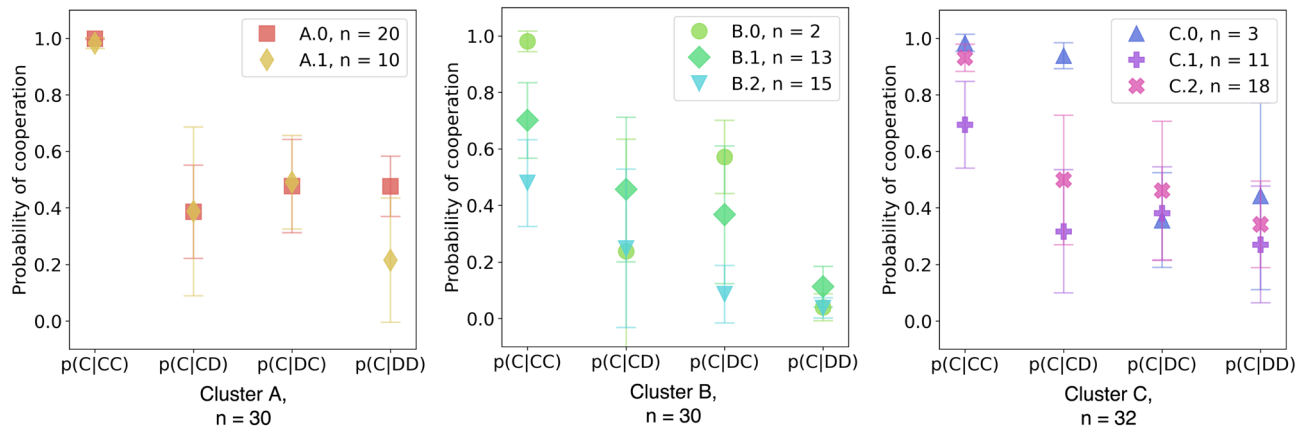


Figure 4. Probability of cooperation given each context per cluster for the FP treatment. In each plot, the conditional probability of cooperating for each context is given. The bars represent the binomial error. Each contextual cluster is divided into sub-clusters that were identified using the frequency of cooperation given a context.

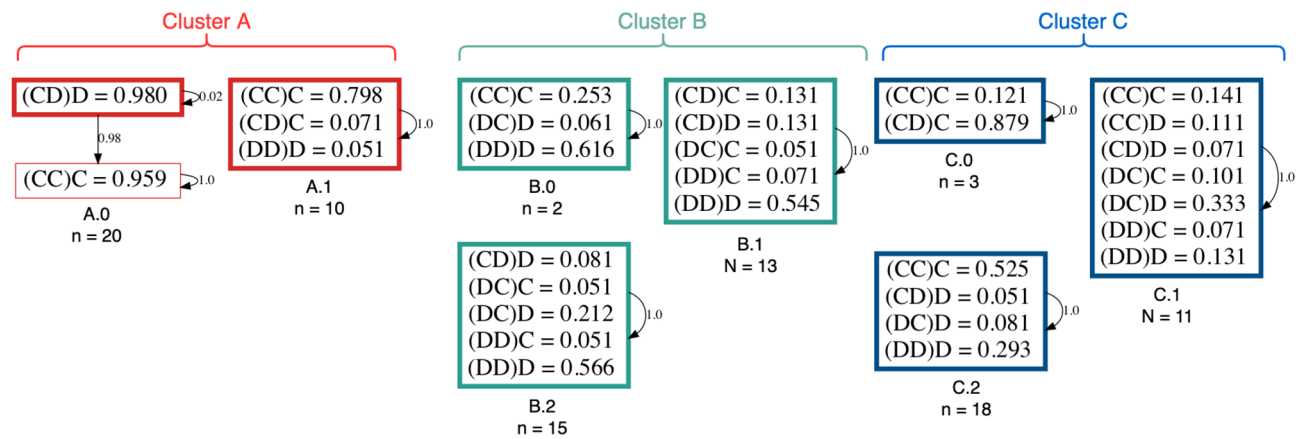


Figure 5. Hidden Markov Models for the FP treatment. Here, the eight sub-clusters found in FP have a HMM that describes each sub-cluster’s strategy for the data over all rounds. Bold rectangles represent the initial state, while the others represent subsequent hidden states. Symbols with a probability lower than 0.05 are not shown, as well as the transition probability between states lower or equal to 0.01.

Although cluster B is divided into three sub-clusters, there are actually two essential ones in terms of the number of participants they represent (see both Fig. 4 and the green HMM in Fig. 5). This next FP cluster mainly experienced mutual defection, and it is situated at the opposite of the behaviors present in the other two FP clusters (see Fig. 3B). Looking at B.1 and B.2 in both figures, participants in sub-cluster B.1 appear to respond more often with cooperation than those in sub-cluster B.2 (see Fig. 4, center panel), given that the triplets $(CD)C$ and $(DD)C$ occur at a higher frequency and that unconditional defection $(DC)D$ occurs with a high frequency in sub-cluster B.2.

Finally, in cluster C of the FP treatment, participants experienced a mixture of contexts, with a preference for mutually cooperative interactions. Clustering this group and inferring the corresponding HMM still revealed some differences (see again Fig. 4, Cluster C and Fig. 5, blue group). Sub-cluster C.0 presents different behavior, members of this sub-cluster unconditionally cooperated despite their opponent’s defection $((CD)C)$. Sub-cluster C.1 had a mixed strategy of reciprocating their opponent’s previous action, but their rate of exploitation $((CC)D = 0.33$ and $(DC)D = 0.07$) is higher than other sub-clusters. Lastly, the sub-cluster C.2 was mainly matching their opponent’s previous action.

While stranger interactions induce defection, cooperative aspirations remain. In the SP treatment, a larger variety of behaviors can be observed, i.e. 11 in total spread again over 3 contextual clusters. Figures 6 and 7 immediately reveal some interesting information about the participants’ behavior in the three clusters found for the SP treatment. Over all three clusters, i.e. D, E, and F, we can observe, through their associated HMM models, that they have different degrees of mutual defection (DD) , with E containing the most (as also was clear from Fig. 3, bottom row, panel A). Moreover, the HMM models for E reveal a high tendency to defect even when the co-player acts cooperatively, even more than the defecting cluster B in FP. The opposite seems

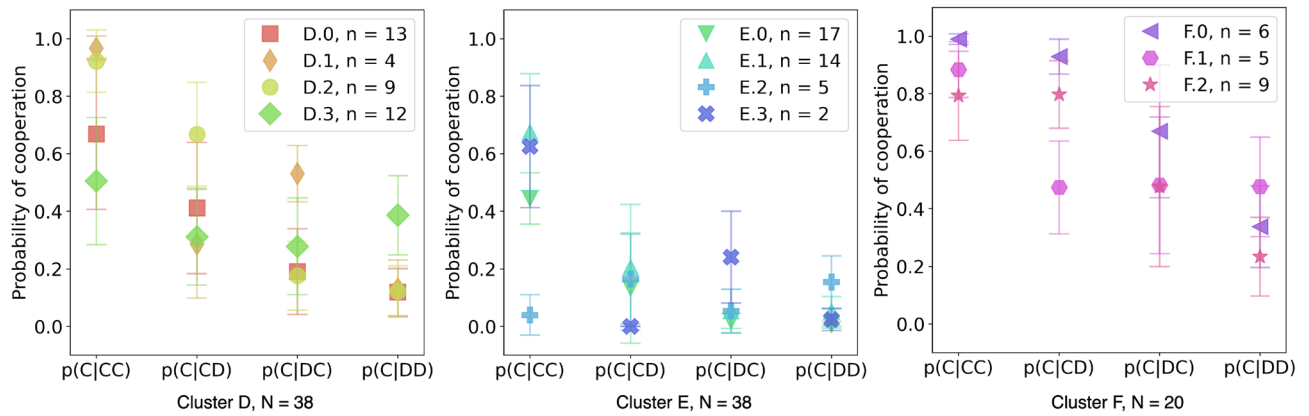


Figure 6. Probability of cooperation given each context for the Shuffled Partners treatment. In the x-axis, it is shown the probability of cooperating given that in the previous round there was a certain context. The bars represent the binomial error. Each cluster is divided into sub-cluster that were divided using the frequency of cooperation given a context.

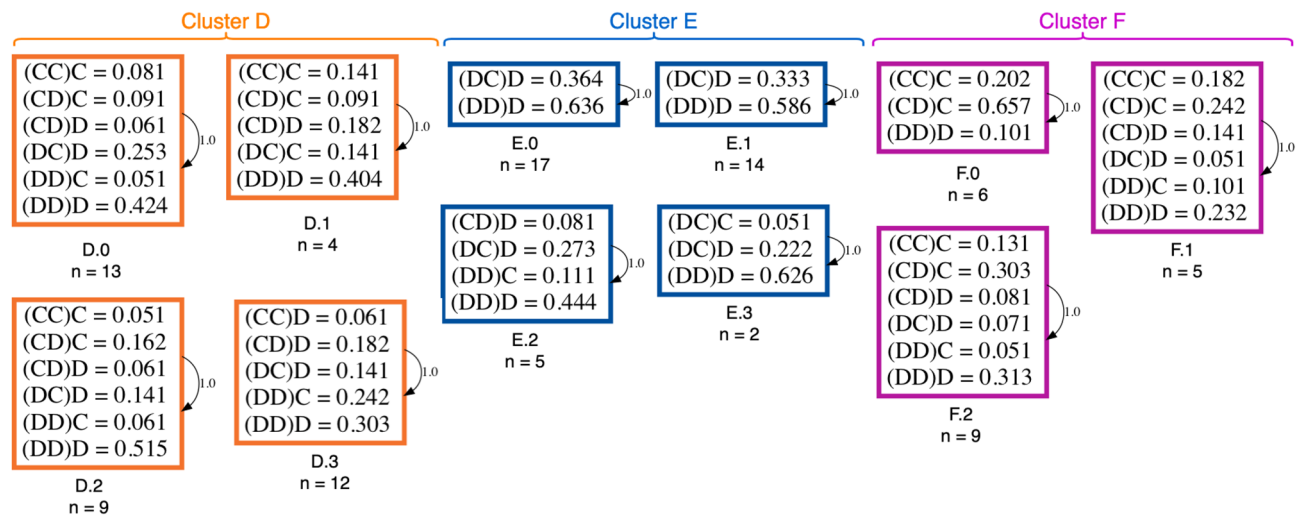


Figure 7. Hidden Markov Models for the Shuffled Partners treatment. Here, the eleven sub-clusters found in SP have a HMM that describes each sub-cluster’s strategy. Symbols with a probability less than 0.05 were not shown, as well as the transition between states less or equal to 0.01.

to occur in cluster F, where mutual cooperation is the highest and players are more likely to ignore a unilateral defection of the co-player.

Cluster D on the other hand, appears to represent intermediate behavior with players still more mutually defecting but the frequency of some other response patterns has increased, which is also visible in Fig. 3B. Sub-cluster D.1 shows to be more conditionally cooperative ((CC)C and (DC)C) than the other sub-clusters in D, i.e. with probability 0.3 (see Fig. 7, cluster D). Players in sub-cluster D.3 and E.2 appear to try to signal their co-players to cooperate given the high frequency of (DD)C in that subgroup to establish cooperation. The cluster D.0 is very similar to the sub-cluster D.2, where the difference is how forgiving and relentless they appear to be in relation to their co-player (see frequency of (CD)C and (DC)D in D.0 versus D.2).

Decision-making simplifies over time and outcomes are decided early. Although the HMM models in Figs. 5 and 7 provide detailed insights into the behaviors observed over a complete (long) IPD experiment, we need to understand whether these behaviors are consistent over time, or whether early decision-making differs from downstream decision-making, which would indicate a learning effect in the experiment. To achieve this goal, the treatment data for each participant is divided into four parts, each consisting of 25 rounds. For each window of 25 rounds, a separate HMM is trained, collectively visualized in the Supplementary Fig. S8 for the FP treatment and Supplementary Fig. S9 for the FP treatment.

For FP one can observe that each strategy seems to converge with time to a dominating decision-making pattern. For example, the differences we observed between sub-clusters A.0 and A.1 are determined by how they act in the first 25 rounds, where it seems participants were still exploring their options. For example, participants in sub-cluster A.1, show a mix of unconditional cooperation ((CD)C), unconditional defection ((DC)D) and

conditional behaviors ((*CD*)*D*, (*DD*)*D* and (*CC*)*C*). After 25 rounds, they coordinated on mutual cooperation for the remainder of the game. Cluster A.0 appears to have led to cooperation thanks to an initial reciprocal behavior ((*CD*)*D*) that was followed by being resolving mistakes if they occur ((*DC*)*C*). Something similar, but then for defection, happens in SP (see Supplementary Fig. S9), where sub-clusters E.0 and E.1 differs essentially in how they react to contexts in the first 25 rounds, yet converge to the same behavioral pattern in the end.

The evolution of sub-cluster B.0 underlines the importance to have long enough experiments: as can be observed, these two participants started out exploring actions in the first 25 rounds, then mutually defected in the next 50 rounds, and finally found a way to cooperate in the last 25. The other two sub-clusters, B.1 and B.2 increase their probability to reciprocate defection ((*DD*)*D*), but participants in sub-cluster B.1 showed from rounds 26-75 a willingness to establish cooperation ((*CD*)*C* = 0.24 and 0.12 respectively) and ended up acting conditionally in the last quarter of the game.

Additionally, although cluster C in FP appeared to have less well-defined behaviors associated with it, one can see that C.0 and C.2 are in fact also very specific in their HMM description. Only for C.1 do we see many different contexts and responses in the emission probabilities of the HMM, yet this abundance of responses remains consistent over all rounds. A very similar situation happens in SP (see Supplementary Fig. S9), where even though in some clusters the strategies became more concise by having less diversity of contexts (see for instance cluster D.0). Yet, the majority of behaviors appear to remain rather stable over time, which is especially clear for cluster E. Overall, the HMM models in the figure reveal explicitly how having new partners at each round induces noise in the decision-process and hinders participants to converge to a clear strategy, relevant for the experiences they have.

FP Participants organize according to their behavioral clusters. So clusters A and B in the FP treatment consist mostly of (forgiving as well as strict) cooperators and defectors respectively. It appears to indicate that behavioral self-selection occurred among the participants in FP. In complex systems and evolutionary game theory, a self-selection mechanism (also called self-organization) occurs when parts of the system appear to reach a stable state⁵⁷. In this case, self-organization brings them to either of the extreme cases, i.e. mostly cooperation or mostly defection. This phenomenon could be happening because of the possibility for imitation of the choices by their neighbors in FP, as argued by Mahmoodi et al.⁵⁸. Members of the third cluster are in some sense still in between, either switching between matching choices or trying to outsmart the co-player. More information about this analysis can be found in the Supplementary information, visualized in Supplementary Fig. S5.

The behavioral self-selection in FP is clearly a consequence of the fixed relationships in that treatment. In that case, previous actions play a more significant role than in the case of the SP treatment. This hypothesis is confirmed by the information obtained in the short questionnaire at the end of the experiment: each participant was asked if her decision was influenced by what the other player did in the previous round (see Supplementary Information for more details). In the SP treatment, 50 players responded that they were influenced (52.08%) by the other's actions. Yet, in the FP treatment, 74 participants responded "yes" to the same question (80.43%). This means that FP participant actions were shaped strongly by actions in the last round, while in SP this shaping did not take place as almost half of the participants did not care about what their opponents did, leaving them with a different approach to decide between cooperating or defect.

Conditional strategies like TFT and WSLS are observed but not consistently used. How do the strategies inferred from the treatment data relate to those proposed in the literature for the origins of cooperative behavior? To answer this question, the HMM results in Supplementary Figs. S8 and S9 need to be considered as the theoretical strategies have been analyzed mostly in fixed partners interactions. Focusing on cluster A in FP first, one can derive that both sub-clusters in A may be associated with TFT-like or a WSLS-like behavior. Considering A.0, one can observe from round 1-25 that this cluster is associated with a reciprocal strategy: It starts out by defecting when the co-player defects (i.e. (*CD*)*D*) but still cooperating when she does ((*DC*)*C*), leading quickly to mutual cooperation ((*CC*)*C*) for the remainder of the FP treatment. Given the strong presence of (*CD*)*D*, this behavior is almost like a form of TFT, around 70% of them started cooperating in the first round as Supplementary Fig. S6 shows, suggesting that participants of this sub-cluster followed the "starting nice" principle as the classical TFT dictates¹⁴.

Behaviors in A.1 also end up in mutual cooperation but achieve this in a different, yet less clear, manner: one could say the participants in A.1 also use a form of reciprocation when the co-player defects but cooperation appears to have been promoted while being generous ((*CD*)*C* in combination with (*CD*)*D*) and signaling to go back to cooperation when both defect ((*DD*)*C*). This behavior is very much associated with the idea of WSLS (Win-Stay, Lose-Shift) or Pavlov as explained by Kraines and Kraines²⁶. When winning ((*DC*) and (*CC*) contexts) continue with the same action, when losing ((*CD*) and (*DD*)) switch. This makes sense in the way that this strategy was designed to promote cooperation in defective environments, and to mimic the stochasticity of biological and social interactions²⁶. Although non-WSLS responses are still present, it appears that together they were sufficient to induce cooperation.

So in the case of fixed partnership interaction, we see successful cases of TFT- and WSLS-like behavior but more often the "implementations" of these strategies did not lead to cooperation, most likely because they were not rigorously applied, as opposed to the case when working with theoretical models. We also see cases of non-conditional behavior, for example in sub-cluster E.0 and E.1 in rounds 26 to 100 where the participants in these clusters defect unconditionally (*DC*)*D* and (*DD*)*D*, which resembles the theoretical strategy AllD. The contrary also holds, where for example, sub-cluster F.0 unconditionally cooperates (*CD*)*C* and (*CC*)*C* during the whole game.

Discussion

As seen in the Introduction, many approaches can be taken to study the strategies in the IPD. In this work, we clustered participants based on the context they experience. This is an important difference from other works since it allows us to see the nuances between people reacting to the same situation. We choose this approach because the act of either cooperating or defecting can mean different things depending on the context of the action. It is not the same to defect to exploit your opponent than to defect to reciprocate your opponents' actions. Moreover, it is not the same to cooperate with a fixed opponent as to cooperate with a stranger, that has been interacting with strangers as well. We used unsupervised methods which make no prior assumptions about what strategies players could use, for this, we tested several methods with the same goal to group participants with similar experiences, and we confirmed there is a significant overlap between different approaches.

A second level clustering (behavioral clustering), on how often they cooperated is necessary to analyze the individual differences given their opponents' actions, and how they signal their intentions in each situation. In fact, in the sub-clusters we found, we could identify players that faced the same context but behaved differently. This means that to understand the strategies people use in the IPD, it is necessary to see what was happening around them in the past and observe how they play in such situations.

Another factor that has an effect on people's strategies in the IPD is their opponent structure. Since we had two treatments, one with fixed co-players and one with shuffled ones, we could see how they act and what strategies they use given each context. Not only players in FP achieved more cooperation, but their strategies looked similar to their co-players. In other words, they were better at self-organizing and the intra-cluster behavior was very similar, as opposed to players of the SP treatment, who had to interact with members of other sub-clusters, who at the same time, had different strategies in mind.

One last hypothesis that we were able to prove here was the learning effect on people's strategies in the IPD. By dividing the analysis into windows of 25 rounds, we could see that the initial quarter was used by the players to explore and try different options, we could even see this effect in a chaotic environment such as SP, where establishing an action plan can be difficult due to its opponent structure.

Using Hidden Markov Models proved to be a handful tool to visualize the strategies and approximate how these underlying factors (in this case, represented as hidden states) shape participants' behavior. Taking a probabilistic approach, rather than a rigid, deterministic one accounts for the stochastic nature of our interactions.

Our findings have implications on how strategies in the IPD are studied: first, performing long experiments allowed us to observe that the participants may not have one but many strategies throughout the game before reaching an equilibrium. This means that when we talk about "how do we really act" we need to account for an exploration stage and a stabilization point. Second, we found that the opponent structure has an effect on people's behavior, and therefore, their strategies. Whether they are more adaptive to their opponents' actions or persistent based on their initial expectations and preferences, this depends on how the game is designed, different strategies might arise. For this reason, we believe that a logical step towards developing these results further is to analyze these strategies in other scenarios, such as N-player IPD, as opposed to the pairwise setting we presented in this paper. Previous work has stated that cooperation could be hindered⁴⁹, but it is uncertain whether they will still be able to reach an equilibrium in their strategies as fast as the pairwise game, and from the strategic point of view, their rationale of what constitutes exploitation or reciprocation changes in this case. Moreover, as we mentioned in the "Methods" section, our methodology can also be applied to more complex strategies such as Memory-2 or higher. Even though subjects may as well be using more sophisticated strategies, the data is not sufficient to really draw conclusions about this. New experiments including a comparative analysis would be required to estimate the presence of such longer-memory strategies, while also assessing their importance relative to the memory-one ones. As this would lead to a paper in itself, this research is beyond the scope of the current work.

Although we could observe some identifiable strategies thanks to the theoretical work behind the IPD, it is still challenging to determine and predict human behavior in these games, and this was represented by the big diversity of options on each HMM. And while taking a probabilistic approach to model decision-making might help us to understand our heuristics, this will continue to be an open issue in research on human cooperation, since all models have their limitations and there are many other scenarios to take into account. Another factor that is well known is the noise that not only our actions produce, but also our perception of what others are doing and their expectations, which they may also have about us. This is a caveat of human interaction to overcome and we are still trying to understand through these experiments.

Data availability

The experimental data and other supplementary information are available at <https://doi.org/10.5061/dryad.37pvmcvmk>. The code to reproduce the results and plots can be found at the Zenodo repository.

Received: 10 January 2022; Accepted: 27 April 2022

Published online: 09 May 2022

References

1. Rand, D. G. & Nowak, M. A. Human cooperation. *Trends Cognit. Sci.* **17**, 413–425. <https://doi.org/10.1016/j.tics.2013.06.003> (2013).
2. Nowak, M. A. Five Rules for the Evolution of Cooperation. *Science* **314**, 1560–1563, <https://doi.org/10.1126/science.1133755> (2006).
3. Gracia-Lázaro, C., Cuesta, J. A., Sánchez, A. & Moreno, Y. Human behavior in Prisoner's dilemma experiments suppresses network reciprocity. *Sci. Rep.* **2**, 325. <https://doi.org/10.1038/srep00325> (2012).
4. Santos, F. C., Santos, M. D. & Pacheco, J. M. Social diversity promotes the emergence of cooperation in public goods games. *Nature* **454**, 213–216, <https://doi.org/10.1038/nature06940> (2008).
5. Perc, M. *et al.* Statistical physics of human cooperation. *Phys. Rep.* **687**, 1–51. <https://doi.org/10.1016/j.physrep.2017.05.004> (2017).

6. Ashlock, D., Ashlock, W. & Umphry, G. An Exploration of differential utility in iterated Prisoner's dilemma. In *2006 IEEE Symposium on Computational Intelligence and Bioinformatics and Computational Biology*, 1–8, <https://doi.org/10.1109/CIBCB.2006.330946> (2006).
7. Dugatkin, L. A. Do guppies play TIT FOR TAT during predator inspection visits? *Behav. Ecol. Sociobiol.* **23**, 395–399. <https://doi.org/10.1007/BF00303714> (1988).
8. Fernández-Domingos, E. *et al.* Delegation to autonomous agents promotes cooperation in collective-risk dilemmas. *arXiv:2103.07710 [cs]* (2021). [ArXiv: 2103.07710](https://arxiv.org/abs/2103.07710).
9. Dawes, R. M. Social dilemmas. *Annu. Rev. Psychol.* **31**, 169–193, <https://doi.org/10.1146/annurev.ps.31.020180.001125> (1980).
10. Lange, P. V., Balliet, D. P., Parks, C. D. & Vugt, M. *Social Dilemmas: Understanding Human Cooperation* (Oxford University Press, 2014).
11. Han, T. A. The Emergence of Commitments and Cooperation. In *Intention Recognition, Commitment and Their Roles in the Evolution of Cooperation: From Artificial Intelligence Techniques to Evolutionary Game Theory Models* (ed. Han, T. A.) 109–121 (Springer, 2013), https://doi.org/10.1007/978-3-642-37512-5_7.
12. Rand, D. G., Ohtsuki, H. & Nowak, M. A. Direct reciprocity with costly punishment: Generous tit-for-tat prevails. *J. Theor. Biol.* **256**, 45–57. <https://doi.org/10.1016/j.jtbi.2008.09.015> (2009).
13. Baek, S. K., Jeong, H.-C., Hilbe, C. & Nowak, M. A. Comparing reactive and memory-one strategies of direct reciprocity. *Sci. Rep.* **6**, 25676, <https://doi.org/10.1038/srep25676> (2016).
14. Axelrod, R. The evolution of strategies in the iterated prisoner's dilemma. In *Genetic Algorithms and Simulated Annealing* 32–41 (Morgan Kaufmann Publishers, 1987).
15. Garcia, J. & van Veelen, M. No strategy can win in the repeated Prisoner's dilemma: Linking game theory and computer simulations. *Front. Robot. AI* **5**, <https://doi.org/10.3389/frobt.2018.00102> (2018).
16. Fudenberg, D. & Maskin, E. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* **54**, 533–554, <https://doi.org/10.2307/1911307> (1986).
17. Axelrod, R. & Hamilton, W. D. The evolution of cooperation. *Science* **211**, 1390–1396, <https://doi.org/10.1126/science.7466396> (1981).
18. Trivers, R. L. The evolution of reciprocal altruism. *Q. Rev. Biol.* **46**, 35–57, <https://doi.org/10.1086/406755> (1971). Publisher: The University of Chicago Press.
19. Nowak, M. Stochastic strategies in the Prisoner's dilemma. *Theor. Popul. Biol.* **38**, 93–112. [https://doi.org/10.1016/0040-5809\(90\)90005-G](https://doi.org/10.1016/0040-5809(90)90005-G) (1990).
20. Reuben, E. & Suetens, S. Revisiting strategic versus non-strategic cooperation. *Exp. Econ.* **15**, 24–43. <https://doi.org/10.1007/s10683-011-9286-4> (2012).
21. Fernández-Domingos, E. *et al.* Timing uncertainty in collective risk dilemmas encourages group reciprocation and polarization. *iScience* **23**, 101752, <https://doi.org/10.1016/j.isci.2020.101752> (2020).
22. Gurven, M. & Winking, J. Collective action in action: Prosocial behavior in and out of the laboratory. *Am. Anthropol.* **110**, 179–190. <https://doi.org/10.1111/j.1548-1433.2008.00024.x> (2008).
23. Wu, J. & Axelrod, R. How to Cope with Noise in the Iterated Prisoner's dilemma. *J. Conflict Resolut.* **39**, 183–189, <https://doi.org/10.1177/0022002795039001008> (1995).
24. Nowak, M. A. & Sigmund, K. Tit for tat in heterogeneous populations. *Nature* **355**, 250–253, <https://doi.org/10.1038/355250a0> (1992).
25. Wedekind, C. & Milinski, M. Human cooperation in the simultaneous and the alternating Prisoner's dilemma: Pavlov versus Generous Tit-for-Tat. In *Proceedings of the National Academy of Sciences* **93**, 2686–2689, <https://doi.org/10.1073/pnas.93.7.2686> (1996).
26. Kraines, D. & Kraines, V. Learning to cooperate with Pavlov an adaptive strategy for the iterated Prisoner's dilemma with noise. *Theory Decis.* **35**, 107–150. <https://doi.org/10.1007/BF01074955> (1993).
27. Martínez-Vaquero, L. A., Cuesta, J. A. & Sánchez, A. Generosity pays in the presence of direct reciprocity: A comprehensive study of 2 × 2 repeated games. *PLOS ONE* **7**, e35135, <https://doi.org/10.1371/journal.pone.0035135> (2012).
28. Milinski, M. TIT for TAT in sticklebacks and the evolution of cooperation. *Nature* **325**, 433–435. <https://doi.org/10.1038/325433a0> (1987).
29. Lombardo, M. P. Mutual restraint in tree swallows: A test of the TIT for TAT model of reciprocity. *Science (New York, N.Y.)* **227**, 1363–1365. <https://doi.org/10.1126/science.227.4692.1363> (1985).
30. Milinski, M. & Wedekind, C. Working memory constrains human cooperation in the Prisoner's dilemma. In *Proceedings of the National Academy of Sciences* **95**, 13755–13758, <https://doi.org/10.1073/pnas.95.23.13755> (1998).
31. Dal Bó, P. & Frechette, G. Strategy choice in the infinitely repeated prisoners' dilemma. Discussion Papers, Research Unit: Economics of Change SP II 2013-311, WZB Berlin Social Science Center (2013).
32. Breiman, L. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Stat. Sci.* **16**, 199–231, <https://doi.org/10.1214/ss/1009213726> (2001).
33. Engle-Warnick, J. Inferring strategies from observed actions: A nonparametric, binary tree classification approach. *J. Econ. Dyn. Control* **27**, 2151–2170. [https://doi.org/10.1016/S0165-1889\(02\)00119-7](https://doi.org/10.1016/S0165-1889(02)00119-7) (2003).
34. Duffy, J. & Engle-Warnick, J. Using Symbolic Regression to Infer Strategies from Experimental Data. In *Evolutionary Computation in Economics and Finance* (ed. Chen, S.-H.) Studies in Fuzziness and Soft Computing, 61–82 (Physica-Verlag HD, 2002), https://doi.org/10.1007/978-3-7908-1784-3_4.
35. Engle-Warnick, J. & Slonim, R. L. Inferring repeated-game strategies from actions: Evidence from trust game experiments. *Econ. Theory* **28**, 603–632. <https://doi.org/10.1007/s00199-005-0633-6> (2006).
36. Engle-Warnick, J. & Slonim, R. L. The evolution of strategies in a repeated trust game. *J. Econ. Behav. Organ.* **55**, 553–573. <https://doi.org/10.1016/j.jebo.2003.11.008> (2004).
37. Engle-Warnick, J. & Ruffle, B. J. The Strategies Behind Their Actions: A New Method to Infer Repeated-Game Strategies and an Application to Buyer Behavior. SSRN Scholarly Paper ID 300500, Social Science Research Network, Rochester, NY (2002). <https://doi.org/10.2139/ssrn.300500>.
38. Kleiman-Weiner, M., Tenenbaum, J. B. & Zhou, P. Non-parametric Bayesian inference of strategies in repeated games. *Econ. J.* **21**, 298–315, <https://doi.org/10.1111/ectj.12112> (2018).
39. Grujić, J., Röhl, T., Semmann, D., Milinski, M. & Traulsen, A. Consistent strategy updating in spatial and non-spatial behavioral experiments does not promote cooperation in social networks. *PLOS ONE* **7**, e47718, <https://doi.org/10.1371/journal.pone.0047718> (2012).
40. Andreoni, J. & Croson, R. Partners versus strangers: Random rematching in public goods experiments. *Handbook Exp. Econ. Results* **1**, 776–783 (2001).
41. Gächter, S. Conditional cooperation: Behavioral regularities from the lab and the field and their policy implications. In *Economics and Psychology: A Promising New Cross-disciplinary Field*, CESifo seminar series, 19–50 (MIT Press, 2007).
42. Heuer, L. & Orland, A. Cooperation in the Prisoner's dilemma: An experimental comparison between pure and mixed strategies. *R. Soc. Open Sci.* **6**, 182142, <https://doi.org/10.1098/rsos.182142> (2019).
43. Dal Bó, P. & Fréchette, G. R. The evolution of cooperation in infinitely repeated games: Experimental evidence. *Am. Econ. Rev.* **101**, 411–429. <https://doi.org/10.1257/aer.101.1.411> (2011).

44. Fudenberg, D., Rand, D. G. & Dreber, A. Slow to anger and fast to forgive: Cooperation in an uncertain world. *Am. Econ. Rev.* **102**, 720–749. <https://doi.org/10.1257/aer.102.2.720> (2012).
45. Fleiß, J. & Leopold-Wildburger, U. Once nice, always nice? Results on factors influencing nice behavior from an iterated Prisoner's dilemma experiment. *Syst. Res. Behav. Sci.* **31**, 327–334. <https://doi.org/10.1002/sres.2194> (2014).
46. Majolo, B. *et al.* Human friendship favours cooperation in the iterated Prisoner's dilemma. *Behaviour* **143**, 1383–1395. <https://doi.org/10.1163/156853906778987506> (2006).
47. Liu, P.-P. Learning about a Reciprocating Opponent in an Iterated Prisoner's Dilemma. In *State University of New York at Stony Brook* (State University of New York at Stony Brook, 2014).
48. Press, W. H. & Dyson, F. J. Iterated Prisoner's dilemma contains strategies that dominate any evolutionary opponent. *Proc. Natl. Acad. Sci. USA* **109**, 10409–10413. <https://doi.org/10.1073/pnas.1206569109> (2012).
49. Grujić, J., Eke, B., Cabrales, A., Cuesta, J. A. & Sánchez, A. Three is a crowd in iterated Prisoner's dilemmas: Experimental evidence on reciprocal behavior. *Sci. Rep.* **2**, 638. <https://doi.org/10.1038/srep00638> (2012).
50. Rabiner, L. A tutorial on hidden Markov models and selected applications in speech recognition. In *Proceedings of the IEEE* **77**, 257–286. <https://doi.org/10.1109/5.18626> (1989).
51. Weiss, R. *et al.* Hmmlern: Unsupervised learning and inference of Hidden Markov Models (2016). <https://github.com/hmmlern/hmmlern>.
52. Maaten, Lvd & Hinton, G. Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**, 2579–2605 (2008).
53. Pedregosa, F. *et al.* Scikit-learn: Machine learning in python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
54. Satopaa, V., Albrecht, J., Irwin, D. & Raghavan, B. Finding a “Kneedle” in a Haystack: Detecting Knee Points in System Behavior. In *2011 31st International Conference on Distributed Computing Systems Workshops*, 166–171. <https://doi.org/10.1109/ICDCSW.2011.20> (2011). ISSN: 2332-5666.
55. Rousseeuw, P. J. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.* **20**, 53–65. [https://doi.org/10.1016/0377-0427\(87\)90125-7](https://doi.org/10.1016/0377-0427(87)90125-7) (1987).
56. Bank, S. G. <https://github.com/xflr6/graphviz> (2021).
57. Yackinous, W. S. Chapter 5-Overview of an Ecological System Dynamics Framework. In *Understanding Complex Ecosystem Dynamics* (ed. Yackinous, W. S.) 83–91 (Academic Press, 2015). <https://doi.org/10.1016/B978-0-12-802031-9.00005-X>.
58. Mahmoodi, K., West, B. J. & Grigolini, P. Self-organizing Complex Networks: Individual versus global rules. *Front. Physiol.* **8**, 478. <https://doi.org/10.3389/fphys.2017.00478> (2017).

Acknowledgements

E.M. and T.L. benefit from the support by the Flemish Government through the AI Research Program and by TAILOR, a project funded by the EU Horizon 2020 research and innovation program under GA No 952215. E.F.D is supported by an F.N.R.S Chargé de Recherche position, grant number 40005955. J.G. was supported by the FWO - Research Foundation Flanders. T.L. is furthermore supported by the F.N.R.S. project with grant number 31257234, the F.W.O. project with grant nr. G.0391.13N, the FuturICT 2.0 (www.futurict2.eu) project funded by the FLAG-ERA JCT 2016 and the Service Public de Wallonie Recherche under grant n° 2010235-ARIAC by DigitalWallonia4.ai.

Author contributions

J.G. designed, performed the experiments, and collected the data. E.M., J.G., T.L., and E.F. analyzed the results. E.M. and J.G. performed the clustering and developed the models. E.M., E.F., and T.L. wrote the manuscript. All authors read and approved the manuscript.

Funding

The funding agency had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-022-11654-2>.

Correspondence and requests for materials should be addressed to E.M.-P.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022