RESEARCH

Supplementary Material: Optimal prediction of decisions and model selection in social dilemmas using block models

Sergio Cobo-López [1], Antonia Godoy-Lorite[2], Jordi Duch[3], Marta Sales-Pardo[1*] and Roger Guimerà[1,4]

*Correspondence:
marta.sales@urv.cat
[1]Departament d'Enginyeria Química, Universitat Rovira i Virgili, Paisos Catalans, 43007 Tarragona, Catalonia, Spain
Full list of author information is available at the end of the article

1 Initial rounds in the empirical data

We observe that the behavior of each player during the first four rounds is erratic, which leads to their behavior being less predictable during those rounds (Fig. S1). After round 4, all rounds are statistically indistinguishable by the metrics discussed in the main text. Therefore, we discard the first four rounds of each player and consider all others as indistinguishable.

2 Single-strategy and multiple-strategy models

In Fig. S2, we illustrate the differences between the single-strategy and the multiple-strategy models.

3 Derivation of the inference equations

3.1 Single-strategy model

We show here the derivation of Eqs. 5 and 6 in the main text, which constitute the basis for the inference in the single-strategy model.

In the single-strategy model, we assume that players and games belong to only one group so that group memberships for player i ($\boldsymbol{\theta}_i$) and game g ($\boldsymbol{\eta}_g$) are binary variables. If σ_i is the group to which player i belongs then $\theta_{ik} = 0 \,\forall k \neq \sigma_i$ and $\theta_{ik} = 1\,, k = \sigma_i$ (and similarly for game memberships η_g). We further assume that the probability of player i cooperating in game g depends exclusively on the group to which they belong $Pr[a_{ig} = C] = p_{\sigma_i \sigma_j}$.

We follow a Bayesian approach to obtain the most plausible group assignments given the observed data (cooperation or defection of players in games) $P(\boldsymbol{\theta}, \boldsymbol{\eta}|A^o)$. To that end, we marginalize over all possible values of the probability of cooperation, $\mathbf{p}$ (Eq. 4 in the main text) as follows:

$$
\begin{aligned}
P(\boldsymbol{\theta}, \boldsymbol{\eta}|A^o) &= \int_{\mathbf{p}} d\mathbf{p} P(\boldsymbol{\theta}, \boldsymbol{\eta}|A^o \mathbf{p}) \\
&= \int_{\mathbf{p}} d\mathbf{p}\, \frac{P(A^o|\boldsymbol{\theta}, \boldsymbol{\eta}, \mathbf{p}) P(\boldsymbol{\theta}, \boldsymbol{\eta}|\mathbf{p}) P(\mathbf{p})}{P(A^0)}
\end{aligned}
\quad , \quad (1)
$$

where we have introduced the Bayes Theorem to express the integrand in terms of the likelihood of the model $P(A^o|\boldsymbol{\theta},\boldsymbol{\eta},\mathbf{p})$, the prior information $P(\boldsymbol{\theta},\boldsymbol{\eta}|\mathbf{p})P(\mathbf{p})$, and the evidence $P(A^0)$, which acts as a normalization constant.

The likelihood of our model is given by:

$$\begin{aligned} P(A^o|\boldsymbol{\theta},\boldsymbol{\eta},\mathbf{p}) &= \prod_{\substack{i\in U\\ g\in\mathcal{C}_i}}\prod_{k\ell} p_{k\ell}\delta_{k\sigma_i}\delta_{\ell\sigma_g}\prod_{\substack{i\in U\\ g\in\mathcal{D}_i}}\prod_{k\ell}(1-p_{k\ell})\delta_{k\sigma_i}\delta_{\ell\sigma_g} \\ &= \prod_{k\ell} p_{k\ell}^{n_{k\ell}^C}(1-p_{k\ell})^{n_{k\ell}^D}\,, \end{aligned} \tag{2}$$

where $\mathcal{C}_i$ and $\mathcal{D}_i$ are the sets of games in which player i cooperates and defects, respectively. $n_{k\ell}^C/n_{k\ell}^D$ are the total number of games in which users in group k cooperate/defect in games belonging to group ℓ.

We assume a flat prior for the $\boldsymbol{\theta}$ and $\mathbf{p}$ and a prior on the $\boldsymbol{\eta}$ that is independent of the $\mathbf{p}$ (see Eq. 3 in the main text) such that

$$P(\boldsymbol{\theta},\boldsymbol{\eta}|\mathbf{p})P(\mathbf{p}) \propto e^{-\alpha F} \qquad \text{with} \qquad F \equiv 1-\sum_{\langle ij\rangle}\boldsymbol{\eta}_i\cdot\boldsymbol{\eta}_j\,, \tag{3}$$

where the sum is over all pairs of neighboring games in the ST plane, so that F is equivalent to the number of neighboring pairs of games (in the ST plane) that don't belong to the same group, and the game aggregation factor α is a positive constant.

Inserting this in the previous equation, we get:

$$P(\boldsymbol{\theta},\boldsymbol{\eta}|A^o) = \frac{e^{-\alpha F}}{\mathcal{Z}}\prod_{k\ell}\int d\mathbf{p}\prod_{k\ell}p_{k\ell}^{n_{k\ell}^C}(1-p_{k\ell})^{n_{k\ell}^D}\,, \tag{4}$$

where $\int d\mathbf{p} \equiv \prod_{k\ell}\int_0^1 dp_{k\ell}$ and $\mathcal{Z}$ is a normalization constant equivalent to the partition function in statistical mechanics.

This integral can be carried out analytically so that,

$$P(\boldsymbol{\theta},\boldsymbol{\eta}|A^o) = \frac{e^{-\alpha F}}{\mathcal{Z}}\prod_{k\ell}\frac{n_{k\ell}^C!\,n_{k\ell}^D!}{(n_{k\ell}^C+n_{k\ell}^D+1)!}\,. \tag{5}$$

By expressing $P(\boldsymbol{\theta},\boldsymbol{\eta}|A^o)$ in terms of the exponential of an energy function $\mathcal{H}$ (Eq. 5 in the main text), we recover the expression for $\mathcal{H}$ in Eq. 6.

3.2 Multiple-strategy model

We want to find the $\boldsymbol{\theta}^*$, $\boldsymbol{\eta}^*$ and $\mathbf{p}^*$ that maximize the posterior $P(\boldsymbol{\theta},\boldsymbol{\eta},\mathbf{p}|A^0)$. We use an Expectation-Maximization algorithm to find the values of those paramaters.

For the mixed-membership model, the logarithm of the likelihood reads:

$$\begin{aligned} \log P(A^o|\boldsymbol{\theta},\boldsymbol{\eta},\mathbf{p}) = &\sum_{i\in U;g\in\mathcal{C}_i}\log\left(\sum_{k\ell}\theta_{ik}p_{k\ell}\eta_{g\ell}\right) \\ &+\sum_{i\in U;g\in\mathcal{D}_i}\log\left(1-\sum_{k\ell}\theta_{ik}p_{k\ell}\eta_{g\ell}\right) \end{aligned} \tag{6}$$

Adding the expression for the exponential prior on the distance between game memberships (Eq. 3) we obtain the following expression for $\log P(\boldsymbol{\theta}, \boldsymbol{\eta}, \mathbf{p}|A^o)$ up to a normalizing constant:

$$\begin{aligned}\log P(\boldsymbol{\theta}, \boldsymbol{\eta}, \mathbf{p}|A^o) &= \sum_{i\in U;g\in\mathcal{C}_i} \log p_{ig} \\ + \sum_{i\in U;g\in\mathcal{D}_i} \log(1-p_{ig}) &- \alpha\left(1-\sum_{\langle gq\rangle} \boldsymbol{\eta}_g\cdot\boldsymbol{\eta}_q\right).\end{aligned} \tag{7}$$

In order to obtain the set of equations that maximize the expression above, we use the following variational trick: we introduce an auxiliary distribution $\omega_{ig}(k,\ell)$, such that $\sum_{k,\ell}\omega_{iq}(k,\ell)=1$, and apply Jensen's inequality $\log\overline{x}\leqslant\overline{\log x}$ so that we can write p_{ig} as:

$$\begin{aligned}\log p_{ig} &= \log\sum_{k,\ell}\theta_{ik}\eta_{g\ell}p_{k\ell} = \log\sum_{k,\ell}\omega_{ig}(k,\ell)\frac{\theta_{ik}\eta_{g\ell}p_{k\ell}}{\omega_{ig}(k,\ell)} \\ &\geqslant \sum_{k,\ell}\omega_{ig}(k,\ell)\log\frac{\theta_{ik}\eta_{g\ell}p_{k\ell}}{\omega_{ig}(k,\ell)},\end{aligned} \tag{8}$$

where the equality holds when $\omega_{ig}(k,\ell)=\frac{\theta_{ik}\eta_{g\ell}p_{k\ell}}{\sum_{k,\ell}\theta_{ik}\eta_{g\ell}p_{k\ell}}$.

Applying the same trick to the other terms we can write

$$\begin{aligned}\log P(\boldsymbol{\theta}, \boldsymbol{\eta}, \mathbf{p}|A^0) &\geqslant \sum_{\substack{i\in U;g\in\mathcal{C}_i\\ k\ell}}\omega^{\mathcal{C}}_{ig}(k,\ell)\log\frac{\theta_{ik}\eta_{g\ell}p_{k\ell}}{\omega_{ig}(k,\ell)} \\ + \sum_{\substack{i\in U;g\in\mathcal{C}_i\\ k\ell}}\omega^{\mathcal{D}}_{ig}(k,\ell)\log\frac{\theta_{ik}\eta_{g\ell}(1-p_{k\ell})}{\omega^{\mathcal{D}}_{ig}(k,\ell)} &- \alpha\left(1-\sum_{\langle gq\rangle}\boldsymbol{\eta}_g\cdot\boldsymbol{\eta}_q\right),\end{aligned} \tag{9}$$

where the equality holds for:

$$w^{C}_{ig}(k,\ell) = \frac{\theta_{ik}\eta_{g\ell}p_{k\ell}}{\sum_{k'\ell'}\theta_{ik'}\eta_{g\ell'}p_{k'\ell'}} \tag{10}$$

$$w^{D}_{ig}(k,\ell) = \frac{\theta_{ik}\eta_{g\ell}(1-p_{k\ell})}{\sum_{k'\ell'}\theta_{ik'}\eta_{q\ell'}(1-p_{k'\ell'})}. \tag{11}$$

The expressions above correspond to the expectation step of the algorithm.

Because we want to obtain the $\boldsymbol{\theta}^\star$, $\boldsymbol{\eta}^\star$,$\mathbf{p}^\star$ that maximize the posterior, we need to maximize the left hand side in Eq. 9. By taking derivatives of the l.h.s in 9 and using Lagrange multipliers for normalization constrains for θ_{ik} and η_{ql}, we obtain the following equations:

$$\theta_{ik} = \frac{\sum_{g\in\mathcal{C}_i}\sum_{\ell} w_{ig}^C(k,\ell) + \sum_{g\in\mathcal{D}_i}\sum_{\ell} w_{ig}^D(k,\ell)}{d_i} \tag{12}$$

$$\begin{aligned} \eta_{g\ell} &= \frac{\sum_{i\in\mathcal{C}_g}\sum_{k} w_{ig}^C(k,\ell) + \sum_{i\in\mathcal{D}_g}\sum_{k} w_{ig}^D(k,\ell)}{d_g + \alpha\sum_{r\in\partial g}\boldsymbol{\eta}_r\cdot\boldsymbol{\eta}_g} \\ &\quad + \frac{\alpha\sum_{r\in\partial g}\eta_{r\ell}\eta_{g\ell}}{d_g + \alpha\sum_{r\in\partial g}\boldsymbol{\eta}_r\cdot\boldsymbol{\eta}_g}, \end{aligned} \tag{13}$$

with $d_i = \sum_{q\in\mathcal{C}_i}\sum_{k\ell} w_{iq}^C(k,\ell) + \sum_{n\in\mathcal{D}_i}\sum_{k\ell} w_{in}^D(k,\ell)$ and $d_q = \sum_{i\in\mathcal{C}_i}\sum_{k\ell} w_{iq}^C(k,\ell) + \sum_{m\in\mathcal{D}_i}\sum_{k\ell} w_{mq}^D(k,\ell)$. $\mathcal{C}_g$ and $\mathcal{D}_g$ are the set of users that cooperate or defect in game g, respectively.

Finally, for $p_{k\ell}$:

$$p_{k\ell} = \frac{\sum_{(i,g)\in\mathcal{C}} w_{ig}^C(k,\ell)}{\sum_{(i,g)\in\mathcal{C}} w_{ig}^C(k,\ell) + \sum_{(m,n)\in\mathcal{D}} w_{mn}^D(k,\ell)}, \tag{14}$$

where $\mathcal{C}$ and $\mathcal{D}$ are the set of (player, game) pairs in which there is cooperation or defection, respectively.

4 Simulated annealing

As discussed in the text, we use simulated annealing to find the group assignments $\boldsymbol{\theta}$ and $\boldsymbol{\eta}$ that maximize the expression for the posterior in Eq. 5 or alternatively minimize the energy $\mathcal{H}_{min}$ in Eq. 6 of the main text. However, solving this problem analytically is computationally infeasible, given the vast number of possible partitions (group assignments) of the system. Therefore, we implement a simulated annealing algorithm to perform this task [1]. The idea of this method is the following: starting from a given partition of players and games whose energy $\mathcal{H}_0$ is known, new partitions are proposed by randomly moving players and games to different groups and energies are computed. The new partitions are automatically accepted if $\mathcal{H}_{new} < \mathcal{H}_0$. Otherwise, they are accepted with probability $e^{-\Delta\mathcal{H}/T}$, where T represents the temperature. In this case, the temperature basically controls the tolerance of the system to switching towards partitions with higher energies. The key point of the simulated annealing is that the temperature gradually decreases with the number of iterations. That way, the system can initially explore the whole landscape and escape from a local minima, for instance. For each value of the temperature, we allow $N_{players}^2 + N_{games}^2$ movements of players and games. Then, we cool the system by a factor $\lambda = 0.99$. That is, $T_{new} = \lambda T_{old}$. Finally, if the system doesn't change its energy after 10 temperature changes, the algorithm automatically stops its execution.

5 Number of groups in the single-strategy and multiple-strategy models

In the single-strategy model, the number of groups is determined automatically by the simulated annealing optimization. Since group plausibilities are calculated by marginalizing exactly over the $\mathbf{p}$ matrices (Eq. 1 above), Eq. 5 above already penalizes complex models, and the optimization will naturally choose the optimal number of groups. The optimal model consists of around 20 groups (depending on the cross-validation fold), although 5 or 6 of them alone typically account for more than 50% of the players.

For the multiple-strategy model, the number of groups needs to be fixed manually. As shown in Fig. S3, we find that the optimal predictions are obtained for K=3 groups of players and L=4 groups of games, although performance is not very sensitive to these values. In fact, for larger values of K and L, the performance is similar but some groups are, in practice, left empty.

6 Game mixtures in the multiple-strategy model

As discussed in the main text, games in the multiple-strategy game belong mostly to a single group, even though, in principle, they could belong to mixtures of groups just as players. In Fig. S4 we show the distribution of entropies for games in the multiple-strategy model for different values of the game aggregation factor α.

7 Robustness of the results

In Figs. 2 and 4 of the main text, we show results for a single fold of the 5-fold cross-validation (except panel **a** in each of them, that shows the average over the five folds). Figures S5 and S6 show equivalent results for a different fold, and are very similar to those in the main text, thus indicating that they are robust.

Figures

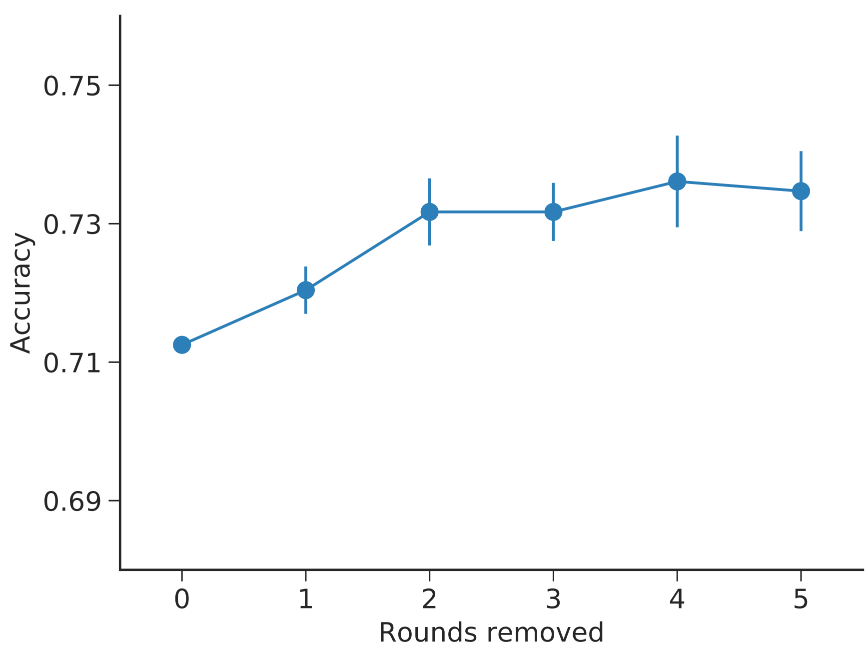

Figure S1 Predictive accuracy of the single-strategy model after removing the first rounds of the players' histories. The error bars show the standard error of the mean of the results for the 5 folds.

Author details
[1]Departament d'Enginyeria Química, Universitat Rovira i Virgili, Paisos Catalans, 43007 Tarragona, Catalonia, Spain. [2]Department of Mathematics, Imperial College London,, , SW7 2AZ London, United Kingdom. [3]Departament de Matemàtiques i Enginyeria Informàtica, Paisos Catalans, 43007 Tarragona, Catalonia, Spain. [4]ICREA, 08010, Passeig Lluís Companys, 08010 Barcelona, Catalonia, Spain.

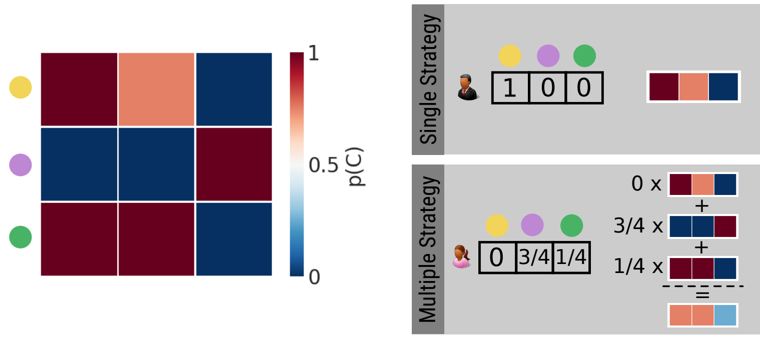

Figure S2 Single-strategy and multiple-strategy models. The matrix of cooperation probabilities (left) indicates the probability with which players in a player group cooperate in games in game group. In the single-strategy model (top right), players belong to a single group, so that their decisions are given directly by the corresponding row in the matrix of cooperation probabilities. In the multiple-strategy model (bottom right), each player can simultaneously belong to different groups of players with given weights. Their decisions are given by the weighted average of the corresponding rows in the matrix of cooperation probabilities.

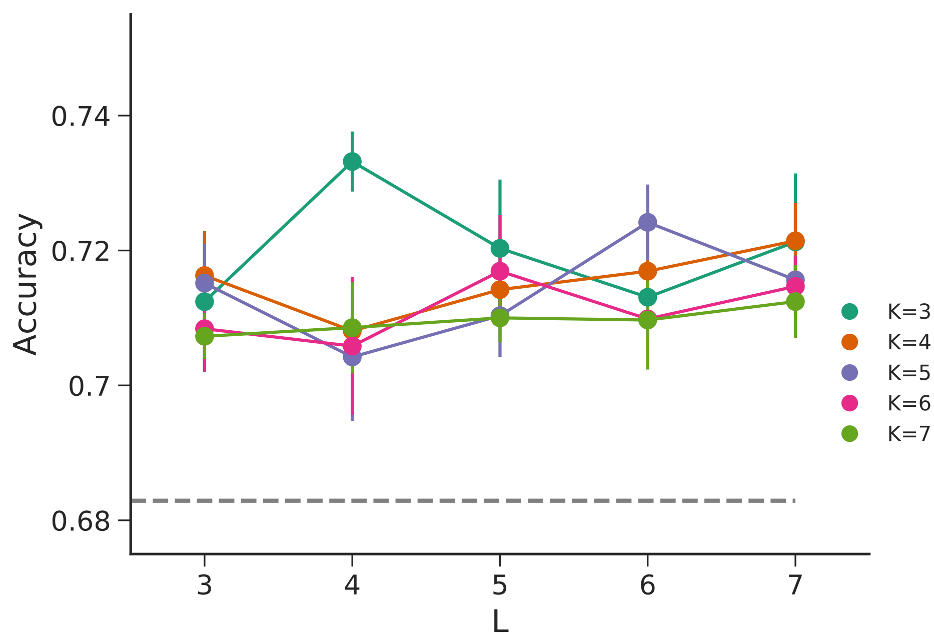

Figure S3 Dependency of the predictive accuracy of the multiple-strategy model for different values of K and L. We show the predictive accuracy for different combinations of K and L (the number of latent groups of players adn games) for $\alpha = 0$. Each point represents the average of a 5-fold cross-validation; error bars indicate the standard error of the mean. Note that for any choice of parameter values the accuracy of the mulitiple-strategy model is above the accuracy of the single-strategy model.

References

1. Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P.: Optimization by simulated annealing. Science **220**, 671–680 (1983)

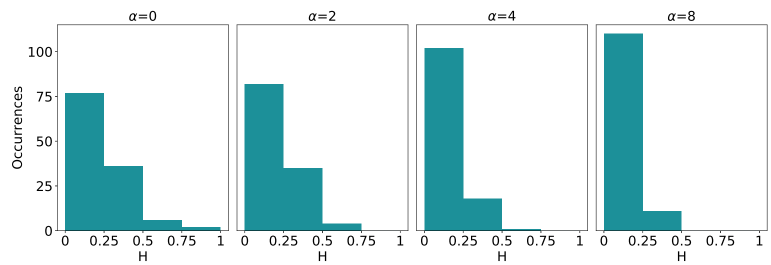

Figure S4 Distribution of Shannon Entropies of the game membership vectors η. Each bin represents the number of occurrences of games significantly belonging to one, two, three or four memberships.

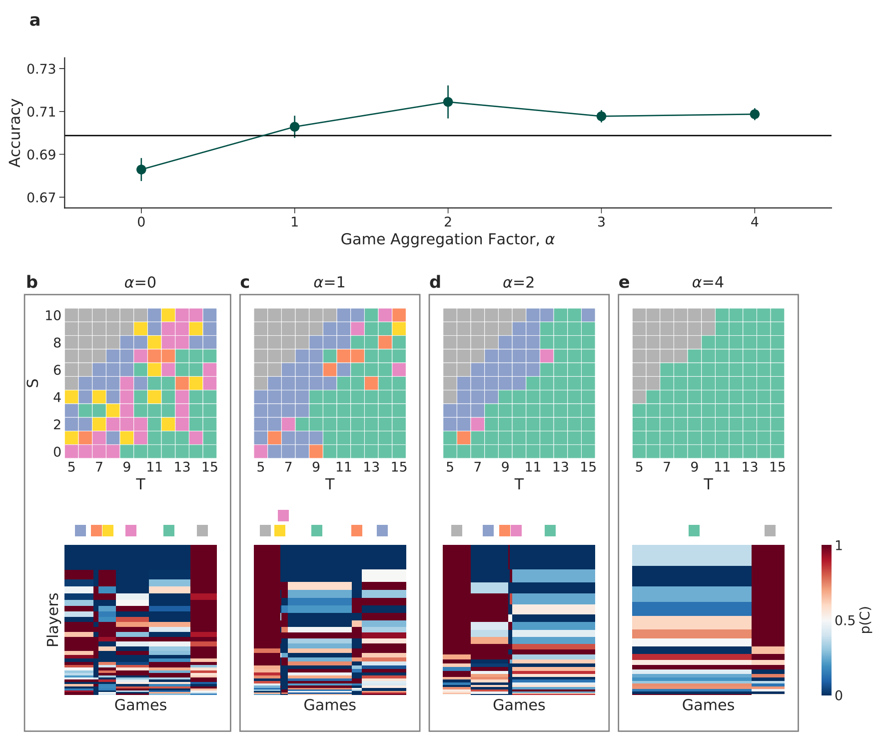

Figure S5 Same as Fig. 2 in the main text for a different split of the 5-fold cross validation. Note how the results we obtain for partitions of games and users are very similar to the ones shown in Fig. 2 of the main text, thus suggesting that the results are robust.

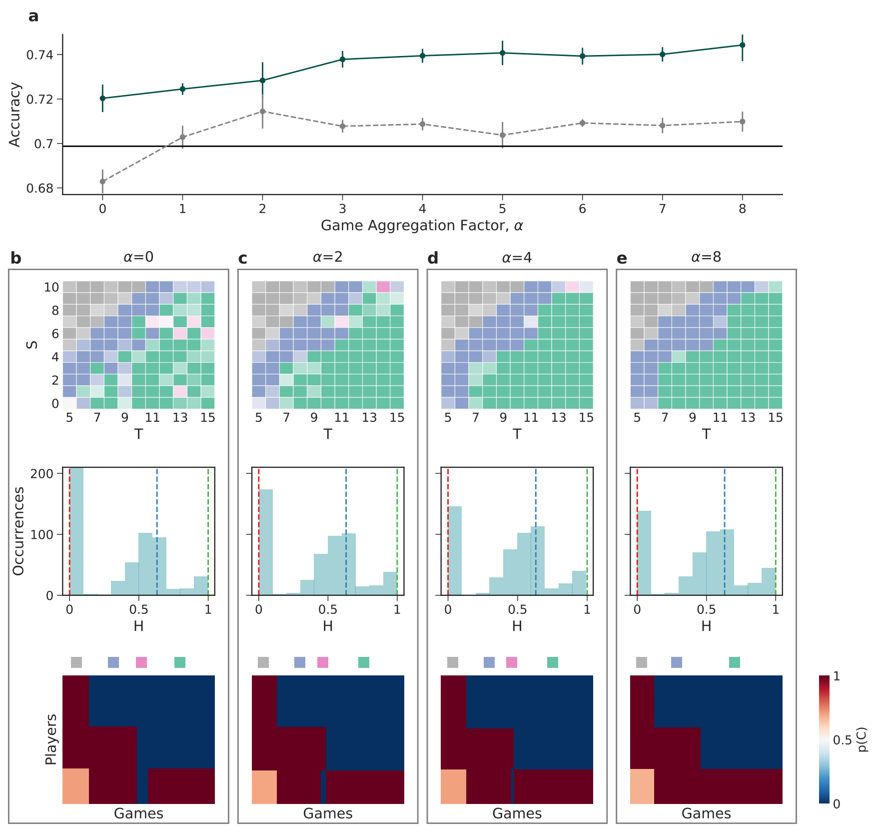

Figure S6 Same as Fig. 4 in the main text for a different split of the 5-fold cross validation (same split as in Fig. S5). Note how rows **b**), **c**) and **d**) are very similar to those shown in Fig. 4 of the main text, thus suggesting that the results are robust.