




Cooperation and deception through stigmergic interactions in human groups

Thomas Bassanetti^{a,b} , Stéphane Cezera^c, Maxime Delacroix^a , Ramón Escobedo^b, Adrien Blanchet^{c,d}, Clément Sire^a , and Guy Theraulaz^{b,d,1} 

Edited by William Bialek, Princeton University, Princeton, NJ; received May 11, 2023; accepted August 30, 2023

Stigmergy is a generic coordination mechanism widely used by animal societies, in which traces left by individuals in a medium guide and stimulate their subsequent actions. In humans, new forms of stigmergic processes have emerged through the development of online services that extensively use the digital traces left by their users. Here, we combine interactive experiments with faithful data-based modeling to investigate how groups of individuals exploit a simple rating system and the resulting traces in an information search task in competitive or noncompetitive conditions. We find that stigmergic interactions can help groups to collectively find the cells with the highest values in a table of hidden numbers. We show that individuals can be classified into three behavioral profiles that differ in their degree of cooperation. Moreover, the competitive situation prompts individuals to give deceptive ratings and reinforces the weight of private information versus social information in their decisions.

social influence | stigmergy | collective intelligence | digital traces | computational modeling

The exchange of social information is the core mechanism by which groups of individuals are able to coordinate their activities and collectively solve problems (1–5). Social information allows individuals to adapt to their environment faster and/or better than through collecting personal information alone (6–10). The use of social information thus provides evolutionary advantages to animal groups and occurs in many contexts, such as foraging, decision-making, division of labor, nest building, or colony defense (1, 2, 11, 12).

Quite often, social information is indirectly shared between individuals: Some of them leave traces of their activities in the environment, and others can use this information to guide their own behavior and inform their own decisions (13). This form of indirect communication, also called stigmergy, in which the trace of an action left on a medium stimulates the performance of a subsequent action which produces another trace and so on, is widely used by animal societies and especially social insects to self-organize their collective behaviors (14–16). These stigmergic interactions that allow the emergence of coordinated activities out of local independent actions likely played a major role in the evolution of cooperativity within groups of organisms (17, 18).

In humans, with the digitalization of society and economies, social information has increasingly taken the form of digital traces, which are the data individuals leave either actively or passively when using the Internet (19–21). New forms of stigmergic processes have been identified since these digital traces are largely exploited in social networks and in electronic commerce, in particular through the use of rating and recommendation systems that can help participants to discover new options and make better choices (22–26). However, individuals do not use social information in the same way. Some individuals exploit it to make their choices, while others may simply ignore it and only use their own private information, or can even go against the message delivered by social information (27). In fact, the same individual can even change the way they provide and use social information depending on the context (28).

Moreover, the use of digital traces is very sensitive to noise and manipulation (29, 30). Indeed, in competitive situations, malicious spammers can manipulate social information by deliberately giving high (respectively, low) ratings to certain low- (respectively, high-) quality items. Therefore, knowing the way individuals share and use digital traces in different contexts is a crucial step to understanding how groups of individuals can cooperate through stigmergic interactions and can exhibit collective intelligence. Despite their increasing importance in human groups, very little research on stigmergic processes has been done so far (31, 32).

The aims of this study are twofold. First, we study through a combination of experiments and computational modeling how indirect interactions between individuals in a human group involving the use of traces allow them to cooperate during an information

Significance

Most online services and applications on the Internet rely on digital traces resulting from choices made by their users, in particular, by means of rating-based recommendation systems. Therefore, it is crucial to understand how such traces affect individual and collective decision-making. We have conducted experiments to measure how groups of individuals interact with digital traces and determine how they could use these traces to cooperate and find the cells with the highest values in a table of hidden numbers. Our experiments and data-driven model show that digital traces spontaneously induce cooperation between individuals. However, the way individuals use the traces to deliver information to others and the reliability of that information largely depends on the degree of competition between individuals.

Author contributions: C.S. and G.T. designed research; T.B., S.C., R.E., A.B., C.S., and G.T. performed research; T.B. and M.D. designed the web interface; T.B. and C.S. designed the model; T.B., C.S., and G.T. analyzed data; and T.B., C.S., and G.T. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2023 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹To whom correspondence may be addressed. E-mail: guy.theraulaz@univ-tlse3.fr.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2307880120/-/DCSupplemental>.

Published October 10, 2023.

search task. Second, we study how a competitive or noncompetitive context influences the way in which individuals exchange and use the social information embedded in the traces of their past actions to perform the information search task.

Through the development of an interactive web application and the use of data-based modeling, we identify the behavioral and cognitive strategies combined with stigmergic interactions that govern individual decisions. The simulation results of our faithful computational model provide clear evidence that the collective behavioral dynamics observed in experiments can be predicted with the precise knowledge of the way individuals use and combine private and social information.

Experimental Design

The experimental setup was designed to investigate in fully controllable conditions how groups of individuals leave and exploit digital traces using a simple five-star rating system similar to the ones used by many online marketplaces and platforms. There, users can evaluate products, services, or sellers and exploit the ratings to help them find the best options corresponding to their expectation.

Here, we study the individual and collective performance of groups of five individuals in a task where each participant has to find the highest values in a 15×15 table of 225 cells, each containing a hidden value (Fig. 1A). In our setup, the cells would represent the available options, and their value would correspond to their intrinsic quality. *SI Appendix, Fig. S1A* presents an example of a table used in our experiments, where the cell values are explicitly shown. Numbers with values ranging from 0 to 99 were randomly distributed in the cells of the table, and *SI Appendix, Fig. S1B* shows the distribution of these cell values. To carry out these experiments, we developed an interactive web application that allows the five group members to independently explore the same table (Fig. 1B and C).

Each experiment includes 20 successive rounds. During each round, each participant has to successively visit and rate three distinct cells. Once a participant discovers the hidden value of a cell, they must rate that cell on a five-star scale. The round ends when everyone in the group has visited and rated three different cells.




Fig. 1. Experimental setup. (A) Screenshot of the table at round $t = 10$, as seen by a participant. In this round, the participant has already visited and rated two cells marked with black crosses. The participant just visited the third cell of value 21 and must rate that cell on a five-star scale. The score of the participant will then depend on the considered rule: In the noncompetitive Rule 1, the score will increase by 0, and by 21 in the competitive Rule 2. (B) Pictures of the experimental room and (C) of the user interface that participants used during the experiment.

At the start of the next round, the color of each cell in the table is updated according to the fraction of stars that have been used to rate the cell since the beginning of the experiment, that is, the number of stars in the cell divided by the total number of stars in all cells. The color scale varies between white (0%) and black (100%) through a gradient of shades of red (*SI Appendix, Fig. S1C*). Thus, the cells that have received the highest fraction of stars since the beginning of the experiment will be clearly visible to all the individuals belonging to the same group. The resulting color map on the table acts like a cumulative long-term collective memory for the group, which is updated at each round. Note that the subjects cannot infer the precise value of the fraction of stars in a cell from its color, but only a rough estimate. However, they can certainly exploit the colors to compare the fraction of stars in the different cells of the table and to identify the cells with a high fraction of stars. Fig. 1A shows an example of a table displaying the participants' ratings as a color map after 10 rounds during one experiment.

We also investigate the impact of a competitive versus a noncompetitive condition on the behavior of participants and the individual and collective performance. In particular, we focus our analysis on the way individuals visit and rate cells and how they use the traces resulting from their ratings and those left by the other group members to guide their choices. In each experiment, we studied the individual and collective behaviors of two groups performing the same task in parallel and independently. In the noncompetitive condition, hereafter called Rule 1, the actions (cell visits and ratings) of the participants do not affect the reward they received at the end of an experiment that always remains constant. On the other hand, in the competitive condition, hereafter called Rule 2, the score of a participant increases at each round by the value of the cells they visit, but remains unaffected by their rating of these cells. Then, the cumulative score of participants over an entire session (12 experiments) determines their monetary reward, which depends on their ranking among the 10 participants and not just among the 5 members of their group (see *Materials and Methods* for the actual payment method).

This experimental design allowed us to study the impact of an intragroup competition since each individual in a group competes with the four other members of their group. However, there is also an implicit intergroup competition since each individual also competes with the five members of the other group for the best rank. *SI Appendix, Fig. S2* illustrates the actions performed by each participant in one group and the color maps associated with the cells in the table resulting from their ratings. *Movies S1A* (Rule 1) and *S2A* (Rule 2) show examples of the dynamics of a typical experimental run where the participants achieved a group score near the observed mean group score. In the corresponding *Movies S3* and *S4*, we present an experimental run where the participants obtained a group score 50% higher than the observed mean group score. *Movie S5* features the same results as *Movies S1–S4* but without the cell values, to better identify the different shades of red and to better reflect what the subjects actually saw during the experiment.

In the next section, we present the results of this experiment mimicking several processes at play in actual five-star rating systems: i) the exploration by the participants of available options (cells in our experiment), which is greatly influenced by their current ratings; ii) the rating on a five-star scale of the options selected by the participants, which ultimately affects the future ratings of these different options. The ratings in our experiments, seen by all participants, are digital traces eliciting stigmergic processes allowing the participants to collectively identify the

best options. In addition, our basic research study also explores the impact of competition in this exploration/rating context, by submitting the participants to noncompetitive or competitive incentives. Although this competitive aspect is less relevant in real-life situations exploiting five-star rating systems, our experimental setup and our modeling approach allow us, more generally, to study the interplay between exploration strategies, rating strategies, private and shared information, and competition.

Results

Collective Dynamics. In this section, we analyze the performance of individuals and groups as well as the dynamics of collective exploration and ratings in both rules. To do so, we introduce a set of precise observables, which are described in detail in *Materials and Methods*: the score of individuals or the mean score of their group; the mean value of the cells weighted by the fraction of stars or the fraction of visits at round t ($p(t)$ and $q(t)$) or up to round t ($P(t)$ and $Q(t)$); the effective number of cells (inverse participation ratio; IPR) over which the stars and visits are distributed at round t and up to round t ; the fidelity F , which quantifies whether the distribution of stars or visits in each cell coincides with the actual distribution of the cell values.

Fig. 2 *A* and *B* show, respectively, the Probability Distribution Function (PDF) of the score S of individuals obtained after the 20 rounds and the score \hat{S} of groups, i.e., the sum of the scores of the individuals belonging to the same group. In Rule 1, all scores are equal to 0. Thus, in order to compare the individual and collective performance in the two rules, each individual is assigned a virtual score computed in the same way as in Rule 2. The mean score is higher in Rule 2, showing that this competitive condition provides a stronger incentive to visit high-value cells than in Rule 1: $\langle S/S_{\max} \rangle = 0.24 \pm 0.01$ in Rule 1 vs. $\langle S/S_{\max} \rangle =$

0.40 ± 0.01 in Rule 2, where $S_{\max} = 5,420$ is the maximum theoretical score.

Fig. 2 *C–F* show that the average value of the visited cells increases with the number of rounds as the participants discover, visit, and rate cells with higher values. Although $p(t)$ and $P(t)$ are higher in Rule 1 than in Rule 2 (Fig. 2 *D* and *F*), the average value of visited cells at round t , $q(t)$, and up to round t , $Q(t)$, are significantly higher in Rule 2 (Fig. 2 *C* and *E*). As we will see later, this apparent paradox is due to the fact that the strategies used by individuals to rate cells in Rule 1 and Rule 2 are very different. In particular, in the competitive Rule 2, some individuals choose to give an average or even a low rating to cells having a high value, presumably to avoid reporting these cells to the other members of their group.

Fig. 2 *G* and *I* show that individuals visit significantly more different cells in Rule 1 than in Rule 2, with $\text{IPR}(\mathbf{Q}(t))$ growing up to the final round $t = 20$ in Rule 1, while it starts decaying after round $t = 7$ in Rule 2. In particular, at the final round $t = 20$ of the experiment, $\text{IPR}(\mathbf{Q}(t = 20))$ is roughly four times larger in Rule 1 than in Rule 2. As we will see in the next section, the lower exploration observed in Rule 2 is mostly due to the fact that the individuals revisit a lot more cells with high values instead of exploring new cells, in order to maximize their score. Moreover, in each round, individuals allocate more stars in Rule 1 compared to Rule 2 (Fig. 2*H*), but overall, they allocate stars to the same number of cells (Fig. 2*J*).

Fig. 2 *K* and *L* show that in both conditions, the fidelity increases with round t , suggesting that the correlations between the participants' visits/ratings and the cell values increase with time. In the final round of the experiment, the fidelity of ratings $F(\mathbf{P}(t = 20), \mathbf{V})$ is significantly higher in Rule 1 than in Rule 2. As we mentioned previously, in Rule 1, the participants explore the table a lot more, and their ratings better reflect the value of the cells that they have discovered.

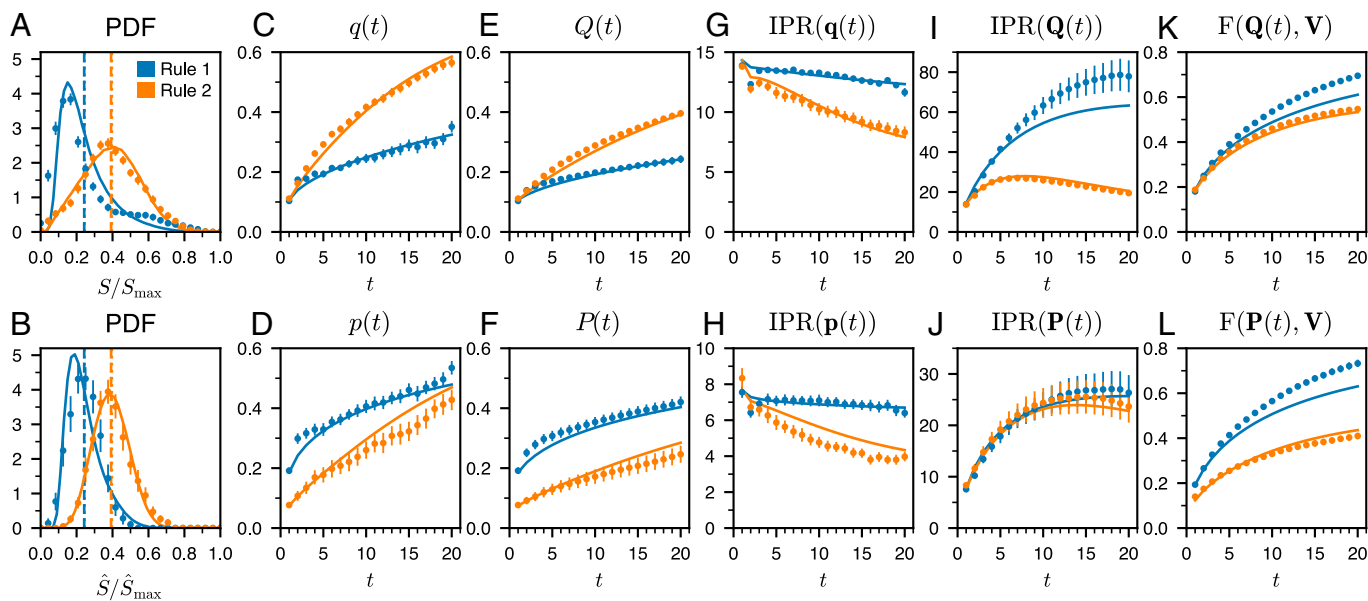


Fig. 2. Collective performance and dynamics of collective exploration and ratings. For the non-competitive Rule 1 (blue) and competitive Rule 2 (orange), the symbols correspond to the experimental results, and the solid lines are the predictions of the model. (A) Probability distribution function (PDF) of the scores of individuals S , and (B) of the groups \hat{S} , respectively normalized by their theoretical maxima S_{\max} and $\hat{S}_{\max} = 5S_{\max}$. The dotted vertical lines are the mean score in the experiment, and the dashed vertical lines are the mean scores in the model. (C) Average value of the cells visited at round t , $q(t)$. (D) Average value of the cells visited weighted by their ratings at round t , $p(t)$. (E) Average value of the cells visited up to round t , $Q(t)$. (F) Average value of the cells visited weighted by their ratings up to round t , $P(t)$. (G) Inverse participation ratio of the visits $\text{IPR}(\mathbf{q}(t))$ and (H) the ratings $\text{IPR}(\mathbf{p}(t))$ at round t , measuring the effective number of cells over which the visits and ratings are distributed at round t . (I) Inverse participation ratio of the visits $\text{IPR}(\mathbf{Q}(t))$ and (J) the ratings $\text{IPR}(\mathbf{P}(t))$ up to round t , measuring the effective number of cells over which the visits and ratings are distributed up to round t . (K) Fidelity to the cell value distribution of the distribution of visits, $F(\mathbf{Q}(t), \mathbf{V})$, and, (L) of ratings, $F(\mathbf{P}(t), \mathbf{V})$.

Individual Behaviors. In this section, we characterize the behaviors of individuals and their strategies to visit and rate cells, i.e., the way they use social information in the form of colored traces resulting from their collective past actions. Moreover, we also quantify the impact of intragroup competition on their behaviors. **Choosing the cells to be visited.** The probabilities of finding the cells with the highest values are higher in Rule 1 than in Rule 2 (Fig. 3 A–C and *SI Appendix*, Fig. S3). In Rule 1, individuals find the best cells more often than would be expected if they had searched randomly, illustrating the cooperative effect induced by the use of the digital traces by individuals within groups. In Rule 2, we observe the opposite phenomenon: Individuals often revisit the cells that they consider high enough to improve their score, without taking the risk of exploring new low-value cells. However, this kind of hedging also hampers their ability to discover even better cells.

We define $V_1(t)$, $V_2(t)$, and $V_3(t)$ as the average of the first-, second-, and third-best values of the cells visited by the participants at round t . Fig. 3 D–F shows that in both conditions, the average values of these 3 cells increase with round t . However,

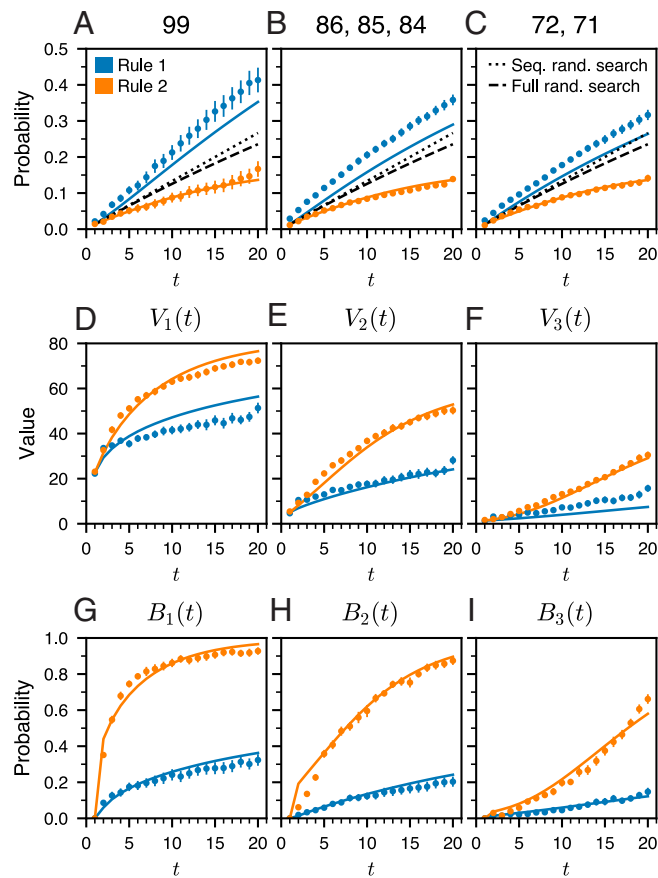


Fig. 3. Quantification of individual behaviors for visiting cells. For the noncompetitive Rule 1 (blue) and the competitive Rule 2 (orange), symbols correspond to the experimental results, while solid lines are the predictions of the model. (A) Probability to find the best cell, of value 99. (B) Probability to find one of the four cells whose values are 86 ($\times 2$), 85, or 84. (C) Probability to find one of the four cells whose values are 72 ($\times 2$) or 71 ($\times 2$). The black dashed and dotted lines correspond to the expected probabilities of two different visiting strategies: *i*) cells chosen randomly (full random search, dashed lines), and *ii*) cells chosen randomly among those that have not been already visited (sequential random search, dotted lines). (D–F) $V_1(t)$, $V_2(t)$, and $V_3(t)$ are, respectively, the value of the first-best cell, second-best cell, and third-best cell visited by the participants, as a function of the round t . (G–I) Probability $B_1(t)$, $B_2(t)$, $B_3(t)$ to revisit the first-best cell, the second-best cell, and the third-best cell of the previous round, as a function of the round $t > 1$.

their average values are higher in Rule 2. Note that this is not in contradiction with the results shown in Fig. 3 A–C. As a matter of fact, in Rule 1, individuals have no incentive to revisit cells with high values, so they continue exploring the table even if they have already found those cells. As already mentioned, in Rule 2, individuals have a clear incentive to revisit cells with high values that they can remember and thus to explore the table less, so that they ultimately discover the cells with the highest values less often.

To confirm this interpretation, we quantify the way individuals revisit cells by defining, for $t > 1$, $B_1(t)$, $B_2(t)$, and $B_3(t)$ as the probability of revisiting at round t the cells with the first-, second-, and third-best values of the previous round ($t - 1$). Fig. 3 G–I show that individuals tend to revisit the cells with the best values, and more so as the value of the visited cells increases over time. In the final round of Rule 2, individuals revisit their first-, second-, and third-best cells of the previous round with respective probabilities 93%, 87%, and 66%. In addition, these observables confirm that individuals explore the table more in Rule 1 than in Rule 2: At any round $t \geq 5$, the values of $B_1(t)$, $B_2(t)$, and $B_3(t)$ in Rule 1 are typically less than one-third of the value in Rule 2.

Altogether, these results illustrate the strong impact of a competitive condition on the way individuals explore the table and select the cells they visit at each round.

Rating the visited cells. *SI Appendix*, Fig. S4 shows the average fraction of stars $\rho(v)$ that has been used to rate cells with value v at the end of the experiment. $\rho(v)$ increases with v , showing that, on average, individuals give higher ratings to cells having high values and also revisit them more often. The experimental data can be fitted to the following functional form:

$$\rho_{\epsilon,\alpha}(v) = \epsilon \frac{1}{N} + (1 - \epsilon) \frac{v^\alpha}{\sum_w N_w w^\alpha}, \quad [1]$$

where $\epsilon \in [0, 1]$ and α are two parameters, $N = 225$ is the total number of cells in a table, and N_v is the number of cells with value v , such that $\sum_v N_v \rho_{\epsilon,\alpha}(v) = 1$. Note that the first term ϵ/N quantifies the fraction of stars uniformly deposited in the cells, while the second term involving α accounts for the fact that high-value cells should attract more stars.

SI Appendix, Fig. S5 shows the average number of stars used to rate a cell as a function of its value v . In Rule 1, the average number of stars increases almost linearly with v . On average, individuals give 1 star to the cells with low values and 4.3 stars to the ones with very high values. In Rule 2, the situation is quite different, individuals give 2.5 stars to low-value cells, and then the average rating decreases to reach a plateau at about 1.5 stars for values higher than $v = 25$. Thus, a cell will receive similar ratings regardless of its value between 35 and 99. This phenomenon suggests that in Rule 2, many participants adopt a noncooperative/deceptive rating strategy, which effectively makes the information conveyed by the digital trace less discriminating. Overall, these results show that individuals give a much fairer rating to the cells they visit in Rule 1, as the examination of the fidelity has previously revealed.

Behavioral profiles of individuals. *SI Appendix*, Figs. S6 and S7 show the average number of stars used to rate a cell as a function of its value v , for each participant, in Rule 1 and Rule 2, leading to three emerging rating patterns. Some individuals rate cells somewhat proportionally to their value, some rate cells independently of their value, and some others give ratings somewhat oppositely proportional to the cell values.

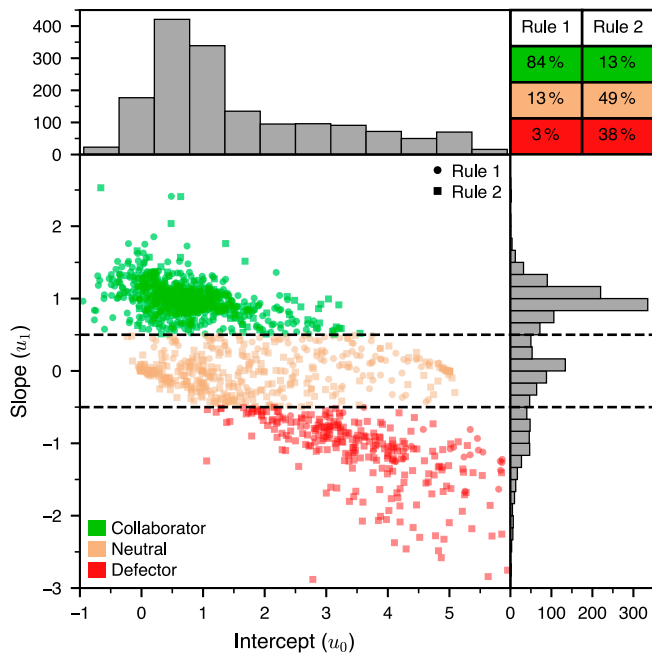


Fig. 4. Behavioral profiles of individuals. (*Bottom-left*) Scatter plot of the values of the two parameters u_0 and u_1 of the linear function, $u_0 + u_1 \times 5v/99$, used to fit each participant's ratings as a function of the value of the visited cells. In the noncompetitive Rule 1, individuals are represented by circles, and in the competitive Rule 2, individuals are represented by squares. The color of the symbols corresponds to the behavioral profile of the individuals: collaborator (green), neutral (brown), and defector (red). The two horizontal lines at $u_{\text{def-neu}} = -0.5$ and $u_{\text{neu-col}} = 0.5$ are the delimitations between the profiles. (*Top-left*) Histogram of the values of u_0 . (*Bottom-right*) Histogram of the values of u_1 . (*Top-right*) The table gives the percentage of individuals for each of the behavioral profiles. See also *SI Appendix, Fig. S8A* (for Rule 1 only) and *B* (for Rule 2 only).

To quantify and classify these three behavioral profiles, we fit the average rating of each individual with a linear function of the cell value v , $u_0 + u_1 \times 5v/99$, where u_0 is the intercept and u_1 is the slope of the line. $u_0 = 0$ and $u_1 = 1$ would correspond to a strict linear rating of cells of value 0 to 99, with 0 to 5 stars. Fig. 4 shows the distribution of u_0 and u_1 for all individuals. We identify three classes of behavioral profiles associated with two thresholds at $u_{\text{def-neu}} = -0.5$ and $u_{\text{neu-col}} = 0.5$ corresponding to the two minima found in the distribution of u_1 . Note that the two thresholds for these three classes are close to the thresholds found using Ward's clustering method on the slope parameter u_1 :

- The ratings of individuals with $u_1 \geq u_{\text{neu-col}}$ increase with the cell values, i.e., they rate cells whose values are the lowest (respectively whose values are the highest) with a small number of stars (respectively a large number of stars; see Fig. 5A). Hereafter, we will dub these individuals as collaborators, since their rating strategy helps the other members of their group to identify the best cells (84% in Rule 1 and 13% in Rule 2).
- Individuals with $u_{\text{def-neu}} \leq u_1 < u_{\text{neu-col}}$ rate cells with almost the same number of stars (on average, three stars in Rule 1, and 1.5 stars in Rule 2) regardless of their values (Fig. 5B). Since the ratings of these individuals do not provide any distinctive information to the other group members, we will dub them as neutrals (13% in Rule 1 and 49% in Rule 2). Note that these neutral individuals do not form a homogeneous group. Indeed, some of them with u_0 close to 0 always give zero or a very few stars whatever the cell value, hence essentially not participating in the rating and the marking of the cells. Some other neutrals

with u_0 close to five always give a large number of stars or even five stars, thus marking all the cells they visit, while others do not have any consistent logic in the way they rate cells. This explains the wide range of intercepts $u_0 \in [0, 5]$ observed for neutrals in Fig. 4. Despite not giving distinctive ratings, most neutrals effectively help the other members of their group to identify the best cells, since they often revisit these cells, and hence make them darker. We also address this point in the section below about optimized agents and in section B.2 of *SI Appendix, Supplementary Text*.

- Individuals with $u_1 < u_{\text{def-neu}}$ rate the cells in the opposite way to collaborators, resulting in deceptive ratings. Indeed, they attribute a small number of stars (respectively a large number of stars) to the cells whose values are the highest (respectively whose values are the lowest; see Fig. 5C). We will call these individuals defectors (3% in Rule 1 and 38% in Rule 2), since we interpret that the strong traces left on cells with very low values are meant to mislead other group members and prevent them from finding the best cells, especially in Rule 2. In addition, they also decide not to share the position of the best cells they have discovered, by giving them low ratings, and hence not marking them on the table.

Fig. 5 A, D, and G show that collaborators mostly rate cells whose values are less than 20 with 1 star, while the cells whose values are greater than 80 are rated with 5 stars. By contrast, Fig. 5 B, E, and H show that for the neutral individuals, the probability

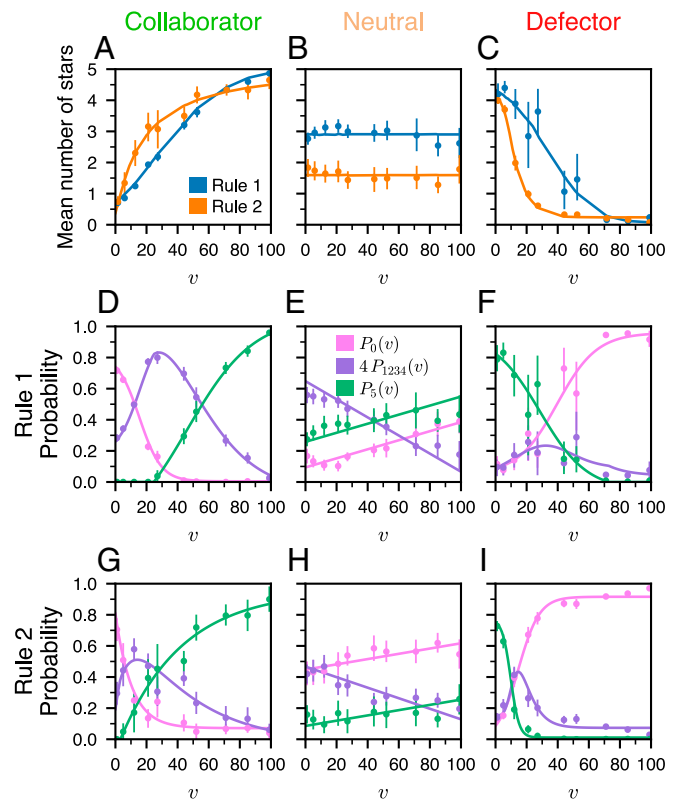


Fig. 5. Rating strategies for the three behavioral profiles. (A–C) Mean number of stars used to rate cells as a function of the cell's value v for (A) collaborators, (B) neutrals, and (C) defectors in the noncompetitive Rule 1 (blue) and the competitive Rule 2 (orange). (D–I) Probability of rating a cell with 0 stars ($P_0(v)$; magenta), 1 to 4 stars ($P_{1234}(v)$; violet) and 5 stars ($P_5(v)$; green) as a function of its value v , for the collaborators, neutrals, and defectors, and for the two rules. The probabilities of rating a cell of value v with 1 to 4 stars have been averaged in $P_{1234}(v)$. The symbols are the experimental data, and the solid lines are the predictions of the model.

of rating a cell with a given number of stars does not depend on the cell value. Finally, Fig. 5 *C*, *F*, and *I* show that the defectors' distribution of ratings presents an inverse pattern compared to that of the collaborators. Defectors poorly rate cells with high values, hence hiding them from the other members of their group. Conversely, they rate cells having low values with a high number of stars, hence misleading others. Ultimately, defectors have access to more information than the other group members. Indeed, the defectors benefit from collaborators who give high ratings to cells having high values. Simultaneously, defectors strategically withhold their knowledge regarding the best cells that they have discovered, by refraining from marking such cells. Thanks to this asymmetric information (33), adopting a defecting behavior can be beneficial in a competitive environment. Indeed, defectors have a higher probability of having the highest score in their group (*SI Appendix*, Fig. S9). However, in the absence of competition, there is no benefit in deception and one should expect fewer defectors. This is what we observe in our experiments, where Fig. 4 (*Inset* table) shows that almost every participant adopts a cooperative behavior in Rule 1, while there is a large fraction of defectors in Rule 2.

Note that the subjects would participate in two experimental runs playing alone before participating in 10 runs with the 4 other members of their group in Rule 1 or Rule 2 (*SI Appendix*, *Supplementary Text*). As expected, when playing alone, the participants behave as collaborators (with themselves), also showing that the participants understood well the principle of the experiment. This was confirmed by asking them to fill an anonymous questionnaire at the end of the session.

Model. We now introduce a stochastic agent-based model to quantitatively identify the strategies for visiting and rating cells and to understand their respective effects on individual and collective performance. In the model, we simulate groups of five agents playing a sequence of 20 consecutive rounds (three visited and rated cells per round), exactly following the actual experimental procedure. The model, described in detail in *Materials and Methods*, consists of two steps that characterize the agents' visit and rating strategies.

The first step accounts for the visit strategy, i.e., which three cells an agent decides to visit in each round. This strategy allows for a variety of behaviors observed in the experiment:

- revisiting the first-, second-, and/or third-best cells already visited in the previous round, depending on their value (private memory; see Fig. 3 *G–I*);
- exploring a marked or unmarked cell (collective memory; see *SI Appendix*, Fig. S4) according to its cumulative fraction of stars represented by the color of the cell in the actual experiment.

The visit strategy is the same for all agents, regardless of their behavioral profile (cooperator, neutral, or defector), as found experimentally, but is allowed to differ for the two conditions, Rule 1 and Rule 2.

The second step of the model addresses the rating strategy, i.e., the number of stars an agent uses to rate a visited cell as a function of its value. In the model, the rating strategy of agents depends on their behavioral profile (Fig. 5 *D–I*) and is different for the two rules.

Model predictions. We consider groups of five agents, hereafter called MIMIC (*Movies S1B* and *S2B*), reproducing the behaviors of human collaborators, neutrals, and defectors. Their behavioral profiles are drawn according to the corresponding fraction

observed in the experiment (*Inset* table of Fig. 4). The parameters for the rating strategies of collaborators, neutrals, and defectors have been estimated by fitting the probability to rate a cell with 0 or 5 stars (see Eqs. 5 and 6 in *Materials and Methods*) to the experimental data (see lines in Fig. 5 *D–I*, and *SI Appendix*, Table S1). As for the parameters for the visit strategy, they have been estimated by minimizing the error between the experimental and the model results for a set of observables, using a Monte Carlo method (*SI Appendix*, Table S2). For all graphs, we ran 1,000,000 simulations, so that the error bars in our simulation results are negligible on the scale of the presented graphs.

Fig. 2 shows that simulations of the model with MIMIC agents quantitatively reproduce the performance of individuals and groups and the observables used to characterize the dynamics of collective exploration and ratings in both rules, as measured in the experiment. The model also quantitatively reproduces the dynamics of the average value of the first-best, second-best, and third-best cells visited by individuals during the different rounds (Fig. 3 *D–F*), along with the probability to revisit each of these three best cells at the next turn (Fig. 3 *G–I*). In addition, the model reproduces fairly the fraction of collaborators, neutrals, or defectors according to their rank at the end of the experiment and the negative impact of the number of defectors on collective performance (*SI Appendix*, Fig. S9). The model also predicts with great accuracy the nontrivial results of Fig. 3 *A–C* and *SI Appendix*, Fig. S3 that were commented above.

These results suggest that the behavioral mechanisms implemented in the model constitute an excellent representation of the processes by which individuals leave and use the traces to guide their choice and how these processes are modulated in the presence of competition between individuals.

Finally, in *SI Appendix*, *Supplementary Text*, we also explore the model predictions for larger group sizes, larger tables, longer durations, and different types of visit and rating strategies.

Optimization of agents' performance according to specific objectives. We have also exploited our model to find agents that are optimized in different situations. To do this, we have used a Monte Carlo method to obtain all the parameters of the model that characterize the corresponding visit and rating strategies.

We first consider a situation in which we wish to maximize the score S (as defined in Rule 2) of five identical agents (Opt-1 agents) in the same group and exploiting the same strategy (*SI Appendix*, Figs. S15 and S19A and Tables S1G and S2). The inspection of the Opt-1 agents' resulting parameters and *SI Appendix*, Fig. S15 show that they essentially only rate cells that have very high values, which they revisit at almost every round so that there is almost no exploration. These Opt-1 agents are strong collaborators, and their average score ($S/S_{\max} = 67\%$) is markedly higher than the score of the human subjects in Rule 2 ($S/S_{\max} = 40\%$). Note that, since the 5 Opt-1 agents are identical, they also maximize the total score of the group. This suggests that a situation where groups would compete (instead of individuals; intergroup instead of intragroup competition) should lead to the emergence of a collaborative behavior within the groups, a situation that we plan to explore experimentally in a future work.

We then consider a situation in which we maximize the score of one agent competing with 4 MIMIC agents (*SI Appendix*, Figs. S16 and S19B and Tables S1H and S2). This scenario represents a more realistic situation where an individual seeks to maximize their score while competing against four other typical individuals. In this condition, the behavior of this optimized agent (Opt-2) is markedly different from that of

Opt-1 agents, since the presence of MIMIC agents behaving as neutrals and defectors forces the Opt-2 agent to adapt its visit and rating strategy to cope with indiscriminate or even false social information. Interestingly, the optimization process leads to a neutral agent assigning 0 star to every visited cell, and hence not participating at all in the rating process. Note that, as already mentioned in the description of neutral agents above (and in section B.2 of *SI Appendix, Supplementary Text*), a neutral agent assigning a nonzero number of stars to visited cells would effectively help the other members of its group to identify the best cells, since it would often revisit these cells. The average score of the Opt-2 agents is $S/S_{\max} = 43\%$, which is only slightly better than the average score of human subjects or MIMIC agents.

However, in our experiment, to obtain the maximum monetary reward, individuals were not strictly required to maximize their score but rather had to optimize their ranking among the 10 individuals in the two groups of five participants. In this condition, the optimized Opt-3 agent competing against 4 (in its group) plus 5 (in the other group) MIMIC agents behaves as a defector (*SI Appendix, Figs. S17 and S19C and Tables S1I and S2*). On average, the Opt-3 agent obtains a rank of 4.57 (compared to a mean rank of 5.5) when ranked among the 10 agents of the two groups and a rank of 2.50 within its own group (mean rank equal to 3). It is remarkable that the model predicts that deception is an emerging behavior in the conditions of our experiment.

Finally, it is interesting to consider the visit and rating strategies maximizing the fidelity of the distribution of ratings to the distribution of cell values in the final round, $F(\mathbf{P}(t = 20), \mathbf{V})$ (*SI Appendix, Fig. S18 and Tables S1 and S2*). If the number of rounds were infinite, the optimal strategy for these agents (Opt-4) would be to explore the table randomly and to rate cells proportionally to their value on a full scale of 0 to 5 stars (corresponding to $u_0 = 0$ and $u_1 = 1$ in Fig. 4). By using this strategy, the agents achieve a fidelity of 0.76 at round 20 (compared to 0.4 in Fig. 2L), and the fidelity would ultimately converge to 1 in the limit of an infinite number of rounds. Clearly, these Opt-4 agents achieve a very mediocre mean score of $S/S_{\max} = 11\%$ compared to that of the previous optimized agents, and even compared to MIMIC agents reproducing the experimental results, and to the human participants. It is worth noting that there could exist a better strategy to maximize the fidelity at round $t = 20$, specifically tailored for the finite 20-round setting used in the actual experiment.

Discussion

The ability to exploit the traces left in the environment by the action of organisms is one of the simplest and oldest mechanisms used to coordinate collective behaviors in biological systems (34–36). In humans, over the past thirty years, the massive development of the Internet, together with applications that extensively use digital traces left voluntarily or not by their users, has reinforced the need to understand how these traces influence individual and collective behaviors (25, 37–39).

In this work, we have measured and modeled the way groups of individuals leave and use digital traces in an information search task implementing a five-star rating system similar to the ones used by many online marketplaces and platforms such as Amazon, TripAdvisor, or eBay, in which users can evaluate products, services, or sellers. Although we certainly do not claim that our experimental setup captures all the processes at play in

these real-life situations, it shares with them an exploration of the available options (cells in our experiment; products for an online store) greatly influenced by their current ratings, and a rating of the selected options by the participants, allowing the ratings to evolve dynamically. However, real rating systems usually provide the users with not only the mean rating of an available option but also the number of ratings for this option, which allows them to modulate their confidence in the different ratings.

Our experiment considered two different rules, with Rule 2 implementing a monetary incentive for participants to perform well, resulting in an explicit competition, absent in Rule 1.

Our experimental results show that groups of individuals can use colored traces resulting from their ratings to coordinate their search and collectively find the cells with the highest values in a table of hidden numbers. These traces constitute a form of long-term collective memory of the past actions performed by the group (21, 40). Combined with the individual short-term memory of the value of the cells already visited, these traces determine the choice of the cells ultimately visited by the participants.

However, our results have also revealed profound disparities in the way individuals use social information resulting from these colored traces to guide them in their tasks and also in the way they choose to deliver information to other group members through their ratings. We have identified three behavioral profiles (collaborators, defectors, and neutrals) that essentially account for the way in which individuals rate cells. Collaborators cooperate by leaving a trace whose intensity positively correlates with the hidden value of the cells, while defectors adopt an opposite behavior. Neutral individuals constitute a sizable fraction of the group members (13% in Rule 1 and 49% in Rule 2) and their ratings are essentially uncorrelated with the actual value of the cells. Yet, the marks that they leave, even if they do not directly inform about the value of the cells, nevertheless induce a cooperative behavior, since neutrals often revisit the high-value cells in a way statistically indistinguishable from the collaborators and defectors.

The information contained in the traces can thus be manipulated by individuals depending on the context, competitive or not, in which the task is performed. Therefore, one may expect that when a situation becomes competitive, individuals should pay less attention to the socially generated traces since the reliability of the information contained in the trace decreases. Previous works in social decision-making have indeed shown that there exists a causal link between mistrust and a decrease in information sharing and that the fear of being exploited can be a reason for group members to withhold accurate information (41, 42). This clearly occurs in Rule 2, where 87% of individuals provide indiscriminate (neutrals) or false (defectors) information, whereas 84% of individuals (collaborators) provide reliable information in Rule 1.

Despite participants achieving higher scores in the competitive Rule 2 than in Rule 1, by exploring less and revisiting the best cells more, the fidelity of the cumulative trace resulting from their ratings is more faithful to the actual distribution of cell values in Rule 1 than in Rule 2. In other words, there is a better relation (more faithful) between the final rating of a cell and its true value in Rule 1 than in Rule 2, although this relation that we measured remains nonlinear.

We used these experimental observations to build and calibrate a model that quantitatively reproduces the dynamics of collective exploration and ratings, as well as the individual and collective performances observed in both experimental conditions.

In particular, this agreement between the model and the experiment is quantified by exploiting a series of subtle observables: PDF of the score, fidelity, inverse participation ratio, probability of revisiting cells depending on their values, etc. Note that an important added value of our model is to offer (via the analysis of its parameters) a direct and quantitative interpretation of the visit and rating strategies for the three observed behavioral profiles of human participants and also for different types of optimized agents. The analysis of individual behaviors combined with the simulations of the computational model shows that competition reinforces the weight of private information (i.e., the individual's memory of the cells already visited) compared to social information (i.e., the collective memory of the group shown on the shared colored table) in the choice of cells that are visited.

The analysis of the model shows that a cooperative effect induced by the trace emerges as soon as there exists a minimal level of marking on cells and that the fidelity of the ratings increases with cooperation. The model also shows that the trace induces weak cooperation even in groups of defectors, provided they rate cells with a large enough number of stars, simply because they revisit the cells whose values are the highest. In this case, individual memory plays a major role in the collective performance of these defectors. Furthermore, the model predicts that the cooperative effect induced by the traces and the average performance of individuals increases with group size. This property results from the stigmergic interactions between individuals that make it possible to amplify at the group level the information about the location of cells whose values are the highest. Similar properties are observed in many species of ants that use pheromone trail laying to coordinate collective foraging activities and to find the best food sources in their environment (43, 44). The model also allowed us to explore the dynamics of the system in different conditions (number of agents and their behavioral strategy, size of the table, number of rounds, etc.) and to investigate the optimal agents' strategy depending on diverse specified objectives. Our analysis shows that the maximal score is obtained for collaborative agents (Opt-1), suggesting that inner-group collaboration should emerge from intergroup competition. Interestingly, the model also predicts that a defector behavior emerges for an agent (Opt-3) aiming at optimizing its rank among the 10 participants of 2 groups, in the same conditions as in our experiment.

As our model was deliberately designed to prioritize relative simplicity, it consequently presents a notable limitation by not incorporating a possible explicit time dependence in the parameters that quantify the visit and rating strategies. Indeed, the perceived importance of a cell with a given color may vary between the beginning and the end of an experimental run. In fact, in the model, the time dependence of a subject's actions only results from the explicit time dependence of the cell colors and of their three best-discovered cells. Again, we did not consider, say, time-dependent visit parameters (ϵ and α parameters), for the sake of simplicity of the model, but also due to the fact that identifying the possible time dependence of these parameters with reasonable statistical accuracy would require a much larger dataset. Yet, despite the model's imperfection in reproducing certain observables, the worst agreement between experimental and model results typically remains within 2 experimental standard errors (for instance, see Fig. 2I for Rule 1). Considering the number and diversity of observables that we have considered (see figures in the main text and *SI Appendix*), this level of agreement can be regarded as very satisfactory, suggesting that the

model grasps the main ingredients of the actual visit and rating dynamics.

Finally, we would like to strongly emphasize that our experimental setup coupled to our predictive model is extremely rich and versatile. Indeed, it can be straightforwardly adapted to the investigation of many other interesting aspects of stigmergic processes as well as the respective impacts of intragroup and intergroup competition on the emergence of cooperation in human groups. In fact, our web application also permits the inclusion of bots (for instance, MIMIC or OPT agents) competing with human subjects in the same group of controllable size, which offers the possibility to investigate the behavior of a subject depending on the composition of their group. Moreover, we have also designed an identical version of our interactive web application which can be deployed on the Internet, and which could be used to conduct large-scale experiments. We plan to explore these different avenues in future works.

Ultimately, understanding and modeling the processes that govern the influence of social information embedded in digital traces on individual and collective behavior is a crucial step to developing personalized decision-making algorithms as well as artificial collective intelligence systems based on stigmergy (26, 45, 46).

Materials and Methods

Ethics Statement. The aims and procedures of the experiments were approved by the Ethics Committee of the Toulouse School of Economics (TSE). All participants provided written consent for their participation.

Experimental Procedure. We conducted two series of experiments, the first one in December 2021 to study the competitive condition (Rule 2) and the second one in December 2022 to study the noncompetitive condition (Rule 1). A total of 175 participants were recruited, of which 75 (40 females, 35 males) participated in experiments with Rule 1 and 100 (47 females, 53 males) participated in experiments with Rule 2. Each participant could participate in a maximum of two different sessions. The participants were mostly students at the University of Toulouse, with an average age of 22.

All experiments were carried out at the TSE Experimental Laboratory. After entering the experimental room and before starting the experiment, the participants signed the consent form; were explained the rules, the payment conditions, and the anonymity warranty; and were asked to shut down their mobile phones. The participants would then be seated in randomly assigned cubicles (anonymously linked to an ID in our database) that prevented interactions between them (Fig. 1B).

Experiments were conducted using a custom-made interactive web application developed in part in collaboration with the company Andil (www.andil.fr). Participants were presented with the same 15×15 table of 225 cells on their respective computer screen, with each cell associated with a hidden value in the range 0–99. Examples of such tables were provided during the instruction phase. The tables used in the experiments were generated by randomly shuffling the same set of values (*SI Appendix, Fig. S1B*). Thus, all tables contained the same set of values, only randomly arranged in the table (*SI Appendix, Fig. S1A*).

We conducted a total of 10 sessions with Rule 1 and 15 sessions with Rule 2. At the beginning of each session, each participant performed two consecutive experiments alone (*SI Appendix, Supplementary Text* for the analysis of these experiments). The main goal was to ensure that each participant understood the use of the web interface and to measure their spontaneous behavior when the only information available was the digital trace resulting from its own activity. Then, the participants were randomly divided into two groups of five and performed 10 successive experiments. During each experiment, the two groups explored different tables that changed during the different experiments.

Each experiment consisted of 20 consecutive rounds, in which each participant had to visit and rate three different cells within a recommended time of 20 s per

round, beyond which a warning would flash on the screen of late participants. A round would end when all participants in the group had visited and rated 3 cells, and the color of the cells in the table would be updated according to a palette of shades of red that translate the fraction of stars allocated to each cell since the start of the experiment (SI Appendix, Fig. S1C). Participants would then move on to the next round.

In the noncompetitive condition (Rule 1), each participant had to find the cells with the highest values in the table, but their actions (visiting and rating cells) were not translated into a score. In the competitive condition (Rule 2), the score of each participant would increase at each round by the value of the three cells they had visited, but it remained independent of the ratings given to these visited cells. Hence, in Rule 2, the participants' main task was to discover the cells with the highest values, while maximizing their score, and ultimately, their payment at the end of the session. Note that we have introduced a notion of score in Rule 1, to compare the results in the two rules (Fig. 2A and B), although, again, the participants in Rule 1 experiments were never told about any notion of score.

Accordingly, all participants were paid the same 10 € at the end of a Rule 1 session. In Rule 2, the 10 participants, from the 2 groups of 5, were ultimately ranked and paid according to their cumulated score at the end of the session. The participant ranked first was paid 20 €, the second and third were paid 15 €, and the participants ranked from the 4th to the 10th place were paid 10 € each.

Observables Used to Quantify the Collective Behavior. We define $p_c(t)$ as the fraction of stars received by a cell c at round t . The set of $p_c(t)$ for all cells c forms a vector $\mathbf{p}(t)$ of size 225 (vectors are shown in boldface). Another vector of interest is the vector $\mathbf{P}(t)$ of the cumulated fraction of stars $P_c(t)$ that have been attributed to each cell from the beginning up to round t included. Similarly, $\mathbf{q}(t)$ and $\mathbf{Q}(t)$ are vectors whose coordinates $q_c(t)$ and $Q_c(t)$ represent the fraction of visits received by each cell at round t and up to round t , respectively.

From the definition of $p_c(t)$ and $P_c(t)$, we can define the average value of cells visited by the participants weighted by their ratings (fraction of stars) at round t , $\rho(t) = \sum_c p_c(t) V_c / v_{\max_1}$, where $v_{\max_1} = 99$ is the highest value of a cell. In general, we have $\rho(t) \leq 1$, and $\rho(t) = 1$ would correspond to all members of a group only giving a nonzero number of stars to the cell of value 99 at round t . Similarly, we define the cumulated quantity, $P(t) = \sum_c P_c(t) V_c / v_{\max_1}$, the average value of cells visited by the participants weighted by their ratings (fraction of stars) up to round t . Hence, $\rho(t)$ and $P(t)$ quantify the instantaneous and cumulated distribution of stars in relation to the value of the visited cells. In particular, a high value of $P(t)$ (in particular, at the final round $t = 20$) indicates that the participants have concentrated the allocation of stars on high-value cells. Conversely, a low value of $P(t)$ suggests a degree of deception, with participants allocating a high fraction of stars to low-value cells, as observed for Rule 2 where many participants are defectors.

In both rules, participants were explicitly asked to discover cells having high values. However, in Rule 2, their score would increase by the value of the cells they visit, thus providing an incentive that affects the way they visit and/or revisit cells during successive rounds. To quantify this (re)visiting behavior, we consider the normalized average value of the cells visited at round t , $q(t) = \sum_c q_c(t) V_c \times 3 / (v_{\max_1} + v_{\max_2} + v_{\max_3})$, where \mathbf{V} is the vector of the cell values V_c , and $v_{\max_1}, v_{\max_2}, v_{\max_3}$ are, respectively, the first-best, second-best, and third-best values of the cells in the table. This observable is normalized so that $q(t) = 1$ corresponds to the best theoretical performance, i.e., when every individual would visit the three best cells of the table at round t . Similarly, we introduce $Q(t)$ that cumulates all visits up to round t and which is defined by the same expression replacing $q_c(t)$ by $Q_c(t)$. Note that, in Rule 2, since the score of the participants is increased by the value of their visited cells, $q(t)$ and $Q(t)$ directly quantify the instantaneous and cumulated performance of the group. In Rule 1, the participants had no notion of score, but $q(t)$ and $Q(t)$ allow us to characterize the dynamics of their visits and to compare it with that for Rule 2.

To quantify the exploration behavior of the table by the participants, we introduce the inverse participation ratio (IPR) of the probability vectors $\mathbf{q}(t)$, $\mathbf{Q}(t)$, $\mathbf{p}(t)$, and $\mathbf{P}(t)$. For a given probability distribution $\mathbf{X} = \{X_c\}$, the IPR of \mathbf{X} is defined as $\text{IPR}(\mathbf{X}) = 1 / \sum_c X_c^2$. For the 4 vectors considered here, the IPR measures the effective number of cells on which the visits or the ratings are concentrated, at round t or up to round t . Indeed, if a probability vector \mathbf{X} is

equally distributed over n cells among N , we have $X_c = 1/n$ on these cells, and $\text{IPR}(\mathbf{X}) = 1 / [n \times (1/n)^2] = n$, showing that the IPR measures the effective number of cells over which a probability distribution is spread.

We are also interested in the relationship between the hidden values of the cells in the table and the fraction of visits or ratings that these cells have received up to round t . This relation is quantified by the fidelity F , which is defined as $F(\mathbf{X}, \mathbf{V}) = \sum_c \sqrt{X_c V_c} / \sum_{c'} \sqrt{V_{c'}}$, where \mathbf{X} is $\mathbf{Q}(t)$ or $\mathbf{P}(t)$. The fidelity F takes values in the interval $[0, 1]$ and is equal to 1 if and only if the probability vector \mathbf{X} is proportional to the vector of cell values \mathbf{V} , which then corresponds to a perfect fidelity. Indeed, the fidelity can be seen as the scalar product between the vector of coordinates $\sqrt{X_c}$ (of unit Euclidean norm, since $\sum_{c'} \sqrt{X_{c'}}^2 = \sum_{c'} X_{c'} = 1$) and the normalized vector proportional to $\sqrt{V_c}$. Hence, the fidelity measures how well aligned these two vectors are and is, in fact, related to the Hellinger distance between the two distributions. In the context of a real-life five-star rating system, a high fidelity of the cumulated ratings $\mathbf{P}(t)$ would indicate that the ratings provide a fair representation of the actual value of the different options. Of course, in this context, these intrinsic values of the available options are generally unknown. But our experimental setup provides a simpler context where this relation between the ratings (or the visits) of the different options (the cells, in our experiment) and their intrinsic value (the cell values) can be investigated.

Model. The stochastic agent-based model includes two components: i) the agents' strategy for visiting cells and ii) their strategy for rating the visited cells.

Visit strategy. In the first round ($t = 1$), the agents have no information; therefore, the selection of the three cells is fully random. For the other rounds ($t > 1$), the agents adopt the following strategy. For each cell $i = 1, 2, 3$ to visit, they either choose the i th-best cell visited in the previous round, of value $V_i(t-1)$, with probability $P_i^R(V_i(t-1))$, or explore other cells with probability $1 - P_i^R(V_i(t-1))$, with:

$$P_i^R(V_i(t-1)) = \begin{cases} 0 & \text{if } V_i(t-1) < a_i \\ \frac{V_i(t-1) - a_i}{99} b_i & \text{if } a_i \leq V_i(t-1) < a_i + \frac{99}{b_i} \\ 1 & \text{otherwise} \end{cases} \quad [2]$$

where a_i and $b_i > 0$ are parameters. An agent never replays a cell of value $V_i(t-1) < a_i$ and always replays a cell of value $V_i(t-1) > a_i + 99/b_i$ (when this threshold is less than 99, the maximum value of a cell). Between these two thresholds, the probability to revisit the i th-best cell linearly interpolates between 0 and 1. The functional form in Eq. 2 is rich enough to be able to capture diverse behaviors, while only using two free parameters for each of the three-best cells, and is, in fact, consistent with indirect measurements of these probabilities.

When an agent does not visit one of the three cells visited in the previous round, it explores other cells in the table. This is done by associating to each cell c a probability $P^E(c, t)$ to be selected at round t :

$$P^E(c, t) = \epsilon \frac{1}{N} + (1 - \epsilon) \frac{P_c^\alpha(t-1)}{\sum_{c'} P_{c'}^\alpha(t-1)} \quad [3]$$

where $P_c(t-1)$ is the fraction of stars deposited in cell c up to time $t-1$, and $\epsilon \in [0, 1]$ and $\alpha > 0$ are parameters. If the selected cell is one of the 3 cells visited in the previous round, another one is selected according to Eq. 3. In Eq. 3, the parameter ϵ controls the amount of exploration of unmarked cells compared to the marked ones: The higher the value of ϵ , the more random, the selection, i.e., independent of the cell color. The exponent α controls the selection of a cell among the marked ones. A high value for α would result in a preferential selection of the highly marked cells, while a small value for α would lead to a more homogeneous selection of a cell among the marked ones. The simple functional form in Eq. 3 is inspired by the experimental results of SI Appendix, Fig. S4, which are well fitted by the similar functional form in Eq. 1.

The values of the eight parameters appearing in Eqs. 2 and 3 and characterizing the visit strategy of MIMIC agents in Rule 1 and Rule 2 are reported in SI Appendix, Table S2.

Rating strategy. Looking at the probability of rating a cell with s stars for each profile (SI Appendix, Fig. S10), one notes that, except for the collaborators in Rule 1, individuals mostly rate a cell with 0 or 5 stars, and that the other ratings with 1, 2, 3, or 4 stars are less common and have a comparable probability. Therefore, in the model, the probabilities of rating a cell with 1 to 4 stars are set equal and are obtained by imposing the probabilistic normalization condition $\sum_{s=0}^5 P_s(v) = 1$, for each value of v . In other words, for $s = 1, 2, 3, 4$, we obtain

$$P_s(v) = P_{1234}(v) = \frac{1}{4}(1 - P_0(v) - P_5(v)). \quad [4]$$

For $s = 0$ and $s = 5$, the probability $P_s(v)$ to rate a cell of value v with s stars is given by

$$P_s(v) = \begin{cases} c_s + d_s \tanh\left(\frac{v - e_s}{99} f_s\right) & \text{for collaborators/defectors} \\ c'_s + d'_s \frac{v}{99} & \text{for neutrals} \end{cases}, \quad [5]$$

where c_s , d_s , e_s , f_s , c'_s , and d'_s are parameters which must satisfy the property that, for all values of v , $P_0(v) + P_5(v) \leq 1$.

However, the $P_{1234}(v)$ approximation is not valid for the collaborators in Rule 1, who use the whole rating scale to rate cells proportionally to their values. Therefore, for these collaborators, we write for $s = 1, 2, 3, 4, 5$,

$$P_s(v) = d''_s \exp\left(-\left(\frac{v - e''_s}{99} f''_s\right)^2\right), \quad [6]$$

where d''_s , e''_s , and f''_s are parameters which must satisfy the property that, for all values of v , $\sum_{s=1}^5 P_s(v) \leq 1$. Finally, we set $P_0(v) = 1 - \sum_{s=1}^5 P_s(v)$.

The functional form of Eqs. 5 and 6 are well adapted to fit the corresponding probabilities observed in the experiment (Fig. 5 D-I and SI Appendix, Fig. S10A), while allowing to capture very diverse behaviors. SI Appendix, Table S1 presents the values of the parameters appearing in the fitting functional forms of Eqs. 5 and 6.

Determination of Model Parameters. For the MIMIC agents, the eight parameters of the visit strategy have been determined by minimizing the error between a set of n round-dependent observables, $O_1(t), \dots, O_n(t)$, measured in the experiment (by averaging them over every experiment for each of the two considered rules) and the corresponding set of observables, $\hat{O}_1(t), \dots, \hat{O}_n(t)$, obtained from extensive simulations of the model (averaging over 1,000,000 numerical experiments for each rule). The error is hence defined by

$$\Delta = \sum_{i=1}^n \frac{\sum_{t=1}^{20} (\hat{O}_i(t) - O_i(t))^2}{\sum_{t=1}^{20} O_i^2(t)}. \quad [7]$$

The set of round-dependent observables considered for the computation of this error Δ consists in the following quantities: $q(t)$, $Q(t)$, $p(t)$, $P(t)$, $\text{IPR}(\mathbf{q}(t))$, $\text{IPR}(\mathbf{Q}(t))$, $\text{IPR}(\mathbf{p}(t))$, $\text{IPR}(\mathbf{P}(t))$, $F(\mathbf{Q}(t), \mathbf{V})$, $F(\mathbf{P}(t), \mathbf{V})$, $V_1(t)$, $V_2(t)$, $V_3(t)$, $B_1(t)$, $B_2(t)$, and $B_3(t)$. We checked that other sets—in particular, smaller sets—of observables would lead to very comparable results (in particular, in Figs. 2 and 3), fitting some observables slightly better and some others slightly worse,

1. S. Camazine et al., *Self-Organization in Biological Systems* (Princeton University Press, 2001).
2. S. Garnier, J. Gautrais, G. Theraulaz, The biological principles of swarm intelligence. *Swarm Intell.* **1**, 3–31 (2007).
3. I. Couzin, Collective minds. *Nature* **445**, 715 (2007).
4. R. L. Goldstone, T. M. Gureckis, Collective behavior. *Top. Cogn. Sci.* **1**, 412–438 (2009).
5. M. Moussaid, S. Garnier, G. Theraulaz, D. Helbing, Collective information processing and pattern formation in swarms, flocks, and crowds. *Top. Cogn. Sci.* **1**, 469–497 (2009).
6. T. R. Zentall, BGG Jr., *Social Learning: Psychological and Biological Perspectives* (Lawrence Erlbaum Associates, Inc., 1988).
7. C. Heyes, B. Galef, *Social Learning in Animals: The Roots of Culture* (Academic Press, 1996).
8. É. Danchin, L. A. Giraldeau, T. J. Valone, R. H. Wagner, Public information: From nosy neighbors to cultural evolution. *Science* **305**, 487–491 (2004).
9. K. N. Laland, Social learning strategies. *Anim. Learn. Behav.* **32**, 4–14 (2004).
10. J. Duboscq, V. Romano, A. Maclntosh, C. Sueur, Social information transmission in animals: Lessons from studies of diffusion. *Front. Psychol.* **7**, 1147 (2016).

and leading to similar results for the functions characterizing the visit strategy in Eqs. 2 and 3.

To minimize the error in Eq. 7, we have used a Monte Carlo method at zero temperature. At each Monte Carlo step, a small random change is introduced in one of the randomly selected parameters. If the error Δ decreases, the new value of the parameter is accepted; otherwise, the old value of the parameter is conserved. The minimization procedure ends when the error stops decreasing. To account for possible multiple local minima of the error, we started the Monte Carlo simulations from several initial values of the parameters. We kept the final parameters leading to the smallest error. Note that the final parameters obtained in different low-error Monte Carlo runs were found to result in similar functions characterizing the visit strategy in Eqs. 2 and 3.

Finally, to obtain the parameters of the visit and rating strategies of the optimized agents (Opt-1, Opt-2, Opt-3, Opt-4), we have exploited a similar zero-temperature Monte Carlo method as described above. However, instead of minimizing an error, we have maximized the score (Opt-1 and Opt-2) or the ranking (Opt-3) of the agent, or the fidelity $F(\mathbf{P}(t = 20), \mathbf{V})$ in the final round (Opt-4).

Computation of the Error Bars. Error bars for the experimentally measured observables correspond to a level of confidence of 68 % and were determined by exploiting the bootstrap method. Bootstrap is a particular type of Monte Carlo method that evaluates the properties of statistical parameters from an unknown probability distribution by repeated random drawings with replacement from a sample (47). The bootstrap method starts by creating M artificial sets of N experiments by drawing with replacement N experiments among the N original ones. This means that some actual experiments can be drawn more than once in an artificial set, while other experiments may not occur in this set. One can then compute a given observable on every artificial set and obtain its distribution, ultimately leading to confidence intervals (CIs). In our case, the independent experiments are the 10 trials played by a group of 5 individuals. Therefore, we have $N = 20$ experiments for Rule 1, and $N = 15$ experiments for Rule 2, and we used $M = 10,000$ artificial sets to generate bootstrap distributions.

For the numerical simulations of the model, the results correspond to an average over 1,000,000 runs, so that the error bars are negligible on the scale of the presented graphs.

Data, Materials, and Software Availability. Code and experimental data have been deposited in Zenodo (48).

ACKNOWLEDGMENTS. This work was supported by grants from the CNRS Mission for Interdisciplinarity (project SmartCrowd, AMI S2C3) and the CNRS Project 80 | Prime ALTHEA. T.B. was supported by a doctoral fellowship from the CNRS. R.E. was supported by Marie Curie Core Grant Funding (grant no. 655235-SmartMass).

Author affiliations: ^aLaboratoire de Physique Théorique, CNRS, Université Toulouse III-Paul Sabatier, 31062 Toulouse, France; ^bCentre de Recherches sur la Cognition Animale, Centre de Biologie Intégrative, CNRS, Université Toulouse III-Paul Sabatier, 31062 Toulouse, France; ^cToulouse School of Economics, CNRS, 31080 Toulouse, France; and ^dInstitute for Advanced Study in Toulouse, 31080 Toulouse, France

11. D. J. T. Sumpter, *Collective Animal Behavior* (Princeton University Press, 2010).
12. T. D. Seeley, *Honeybee Democracy* (Princeton University Press, 2010).
13. L. Giuggiolli, J. R. Potts, D. I. Rubenstein, S. A. Levin, R. M. May, Stigmergy, collective actions, and animal social spacing. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 16904–16909 (2013).
14. P. P. Grassé, La reconstruction du nid et les coordinations interindividuelles chez *Bellicositermes natalensis* et *Cubitermes* sp. la théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs. *Insectes Soc.* **6**, 41–80 (1959).
15. G. Theraulaz, E. Bonabeau, A brief history of stigmergy. *Artif. Life* **5**, 97–116 (1999).
16. F. Heylighen, Stigmergy as a universal coordination mechanism. II: Varieties and evolution. *Cogn. Syst. Res.* **38**, 50–59 (2016).
17. R. Chiong, M. Kirley, "The evolution of cooperation via stigmergic interactions" in *2012 IEEE Congress on Evolutionary Computation* (Princeton University Press, London, UK, 2012), pp. 1–8.
18. F. Heylighen, Stigmergy as a universal coordination mechanism. I: Definition and components. *Cogn. Syst. Res.* **38**, 4–13 (2016).
19. M. J. Doyle, L. Marsh, Stigmergy 3.0: From ants to economies. *Cogn. Syst. Res.* **21**, 1–6 (2013).

20. F. Cochoy, J. Hagberg, M. McIntyre, *Digitalizing Consumption: How Devices Shape Consumer Culture* (Routledge, 2017).
21. R. K. Baltzersen, *Cultural-Historical Perspectives on Collective Intelligence* (Cambridge University Press, 2022).
22. H. V. D. Parunak, "A survey of environments and mechanisms for human-human stigmergy" in *Environments for Multi-Agent Systems II*, D. Weyns, H. Van Dyke Parunak, F. Michel, Eds. (Springer, Berlin/Heidelberg, Germany, 2006), pp. 163-186.
23. L. Lü *et al.*, Recommender systems. *Phys. Rep.* **519**, 1-49 (2012).
24. T. Hennig-Thurau, A. Marchand, P. Marx, Can automated group recommender systems help consumers make better choices? *J. Mark.* **76**, 89-109 (2012).
25. S. A. Golder, M. W. Macy, Digital footprints: Opportunities and challenges for online social research. *Ann. Rev. Sociol.* **40**, 129-152 (2014).
26. M. Jesse, D. Jannach, Digital nudging with recommender systems: Survey and future directions. *Comput. Hum. Behav.* **3**, 100052 (2021).
27. B. Jayles *et al.*, How social information can improve estimation accuracy in human groups. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 12620-12625 (2017).
28. W. Steinel, S. Utz, L. Koning, The good, the bad and the ugly thing to do when sharing information: Revealing, concealing and lying depend on social motivation, distribution and importance of information. *Organ. Behav. Hum. Decis. Process.* **113**, 85-96 (2010).
29. J. L. Herlocker, J. A. Konstan, L. G. Terveen, J. T. Riedl, Evaluating collaborative filtering recommender systems. *ACM Trans. Inf. Syst.* **22**, 5-53 (2004).
30. I. Gunes, C. Kaleli, A. Bilge, H. Polat, Shilling attacks against recommender systems: A comprehensive survey. *Artif. Intell. Rev.* **42**, 767-799 (2014).
31. L. Correia, A. M. Sebastião, P. Santana, On the role of stigmergy in cognition. *Prog. Artif. Intell.* **6**, 79-86 (2017).
32. N. Coucke, M. K. Heinrich, A. Cleeremans, M. Dorigo, HuGoS: A virtual environment for studying collective human behavior from a swarm intelligence perspective. *Swarm Intell.* **15**, 339-376 (2021).
33. S. Balakrishnan, M. P. Koza, Information asymmetry, adverse selection and joint-ventures: Theory and evidence. *J. Econ. Behav. Organ.* **20**, 99-117 (1993).
34. E. S. Gloag *et al.*, Stigmergy: A key driver of self-organization in bacterial biofilms. *Commun. Integr. Biol.* **6**, e27331 (2013).
35. E. S. Gloag, L. Turnbull, C. B. Whitchurch, Bacterial stigmergy: An organising principle of multicellular collective behaviours of bacteria. *Scientifica* **2015**, 1-8 (2015).
36. A. Khuong *et al.*, Stigmergic construction and topochemical information shape ant nest architecture. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 1303-1308 (2016).
37. S. Karanasios *et al.*, Making sense of digital traces: An activity theory driven ontological approach. *J. Am. Soc. Inform. Sci. Technol.* **64**, 2452-2467 (2013).
38. K. K. Loh, R. Kanai, How has the internet reshaped human cognition? *Neuroscientist* **22**, 506-520 (2016).
39. J. Firth *et al.*, The "online brain": How the internet may be changing our cognition. *World Psychiatry* **18**, 119-129 (2019).
40. B. Thierry, G. Theraulaz, J. Y. Gautier, B. Stiegler, Joint memory. *Behav. Proc.* **35**, 127-140 (1996).
41. D. E. Zand, Trust and managerial problem solving. *Adm. Sci. Q.* **17**, 229-239 (1972).
42. W. Steinel, C. K. D. Dreu, Social motives and strategic misrepresentation in social decision making. *J. Pers. Soc. Psychol.* **86**, 419-434 (2004).
43. S. Goss, S. Aron, J. L. Deneubourg, J. M. Pasteels, Self-organized shortcuts in the argentine ant. *Naturwissenschaften* **76**, 579-581 (1989).
44. R. Beckers, J. L. Deneubourg, S. Goss, J. M. Pasteels, Collective decision making through food recruitment. *Insectes Soc.* **37**, 258-267 (1990).
45. F. Heylighen, L. Apostel, Collective intelligence and its implementation on the web: Algorithms to develop a collective mental map. *Comput. Math. Organ. Theory* **5**, 253-280 (1999).
46. W. Zhang, H. Zhao, Y. Jiang, Z. Jin, Stigmergy-based construction of internetware artifacts. *IEEE Softw.* **32**, 58-66 (2015).
47. A. C. Davison, D. V. Hinkley, *Bootstrap Methods and their Application* (Cambridge University Press, 1997).
48. T. Bassanetti *et al.*, Cooperation and deception through stigmergic interactions in human groups. *Zenodo*. <https://doi.org/10.5281/zenodo.7924494>. Deposited 11 May 2023.