



HAL
open science

Impact of different competition schemes on stigmergic cooperation processes in human groups

Thomas Bassanetti

► **To cite this version:**

Thomas Bassanetti. Impact of different competition schemes on stigmergic cooperation processes in human groups. Symbiosis. Université de Toulouse, 2024. English. NNT : 2024TLSES062 . tel-04674306

HAL Id: tel-04674306

<https://theses.hal.science/tel-04674306>

Submitted on 21 Aug 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Doctorat de l'Université de Toulouse

préparé à l'Université Toulouse III - Paul Sabatier

Impact de différents schémas de compétition sur les
processus de coopération stigmergiques au sein de groupes
humains

Thèse présentée et soutenue, le 4 avril 2024 par

Thomas BASSANETTI

École doctorale

SDM - SCIENCES DE LA MATIERE - Toulouse

Spécialité

Physique

Unité de recherche

LPT - Laboratoire de Physique Théorique

Thèse dirigée par

Clément SIRE et Guy THERAULAZ

Composition du jury

M. David CHAVALARIAS, Rapporteur, CNRS

Mme Laëtitia GAUVIN, Rapporteur, IRD

M. Alain BARRAT, Examineur, CNRS

M. Umberto GRANDI, Examineur, Université Toulouse Capitole

M. Clément SIRE, Directeur de thèse, CNRS

M. Guy THERAULAZ, Co-directeur de thèse, CNRS

Impact of Different Competition Schemes on Stigmergic
Cooperation Processes in Human Groups

*À Pascal,
Pour avoir cultivé mon esprit scientifique.
Je sais que tu aurais aimé lire ces pages.*

Psychohistory was the quintessence of sociology; it was the science of human behavior reduced to mathematical equations. The individual human being is unpredictable, but the reactions of human mobs, Seldon found, could be treated statistically.

Isaac Asimov
(Second Foundation, 1953)

Contents

List of Figures	xvii
List of Tables	xix
List of Definitions	xxi
1 Introduction	1
1.1 Introduction	2
1.2 Collective Intelligence in Human Groups	3
1.2.1 Definition	3
1.2.2 Types of collective problem-solving	3
1.2.3 Social information	4
1.2.4 Sharing social information	4
1.2.5 Social influence	5
1.3 Stigmergic Problem-Solving	5
1.3.1 Stigmergy as a universal coordination mechanism	5
1.3.2 Different types of stigmergy	6
1.3.3 Stigmergic processes in bacterial and animal societies	6
1.3.4 Stigmergic processes in human groups	8
1.4 Cooperation and Competition	10
1.4.1 Cooperation in human groups	10
1.4.2 The evolution of cooperation	10
1.4.3 The dual nature of competition	11
1.5 Enhancing Collective Intelligence	12
1.5.1 Exchange of social information	12
1.5.2 Size and composition of the group	13
1.5.3 Network structure	14
1.5.4 The role of noise in group decision-making	14
1.5.5 Social nudges and bots	15
1.6 Outline of the Thesis	15
2 Experimental Design	17
2.1 The STIGMER game	18
2.1.1 Grid and cell values	19
2.1.2 Round	19
2.1.3 Colors	20
2.1.4 Score	23
2.2 Implementation of the Game	24
2.2.1 Web application	24
2.2.2 Visualization application	27

2.3	Experimental Setup	27
2.3.1	Experimental session	27
2.3.2	Payment	30
2.3.3	Data	30
2.3.4	Participants	31
2.3.5	Ethics statement	31
2.4	Summary of the Conducted Experiments	32
3	Main Observables	35
3.1	Notations	36
3.2	Quantifying the Collective Behavior	37
3.2.1	Normalized group score	37
3.2.2	Fractions of visits and fractions of stars	38
3.2.3	Visit and rating performances	39
3.2.4	Inverse participation ratio	40
3.2.5	Fidelity	41
3.3	Quantifying Individual Visit Behaviors	42
3.3.1	Normalized individual score	42
3.3.2	Values of visited cells	43
3.3.3	Probability of revisiting cells from the previous round	44
3.3.4	Probability of finding cells with the highest values	44
3.4	Quantifying Individual Rating Behaviors	45
3.4.1	Mean rating	45
3.4.2	Fitting of the mean rating	46
3.4.3	Fractions of ratings	47
3.5	Computation of the Error Bars: the Bootstrap Method	47
4	Model	51
4.1	General Description of the Model	52
4.2	Visit Strategy	53
4.2.1	Revisiting the cells visited during the previous round	53
4.2.2	Exploration of the grid	54
4.3	Rating Strategy	55
4.3.1	Using the mean rating	55
4.3.2	Direct specification of the distribution functions	55
4.4	Types of Agents	57
4.4.1	Mimic agents	57
4.4.2	Optimized agents	57
4.4.3	Other agents	57
4.5	Determination of Model Parameters	57
4.5.1	Fitting	58
4.5.2	Monte Carlo method	58

5	Impact of Competitive and Non-Competitive Conditions	59
5.1	Introduction	60
5.2	Experimental Design	61
5.3	Collective Behavior	62
5.4	Individual Behavior	66
5.4.1	Visiting cells	67
5.4.2	Rating cells	69
5.4.3	Behavioral profiles of individuals	71
5.5	Model Predictions	77
5.5.1	Modeling human behavior	78
5.5.2	Comparison between theoretical and experimental results	78
5.5.3	Impact of the number of rounds and group size on individual performance and collective dynamics	81
5.5.4	Impact of the rating strategy on agents' performance and the fidelity of ratings	84
5.5.5	Optimization of agents' performance according to specific objectives	86
5.6	Solo Games	92
5.7	Discussion	92
6	Impact of Group Composition	97
6.1	Introduction	98
6.2	Experimental Design	99
6.2.1	The experiments	99
6.2.2	The bots	100
6.3	Impact of Bots on Participant's Performance	102
6.4	Model of the Visit and Rating Strategies	111
6.5	Model for predicting the Human Participants' Behavioral Profiles	114
6.5.1	Cues available to human participants	114
6.5.2	Linear model for predicting individual behavioral profiles	116
6.5.3	Application of the model	117
6.5.4	Interpretation of the model	119
6.6	Discussion	120
7	Impact of Intragroup and Intergroup Competition	123
7.1	Introduction	124
7.2	Experimental Design	124
7.3	Performance and Behavioral Profiles of the Participants	126
7.4	Model for Predicting Individual Behavioral Profiles	130
7.4.1	Three quantities characterizing the payment schemes	130
7.4.2	Application of the model	132
7.5	Discussion	134

8 Conclusion	137
8.1 Summary of Main Results	137
8.1.1 A game for studying stigmergy	138
8.1.2 Cooperation through stigmergic processes	138
8.1.3 A model of human behaviors	139
8.1.4 Impact of the cooperation level of the group	140
8.1.5 Impact of the type of competition	141
8.2 Limitations	142
8.2.1 Non independence of the games in an experimental session	143
8.2.2 Other means to tune the competition	143
8.2.3 Deception is (too) easy	143
8.2.4 Classification of neutrals	143
8.3 Perspectives	144
8.3.1 Evolution of cooperation	144
8.3.2 Knowledge of the performance of both groups	145
8.3.3 Adaptation in dynamic landscapes	145
8.3.4 No private memory	145
8.3.5 Varying group sizes	146
8.3.6 Analysis of the questionnaires	146
 Appendices	 147
 A Documents Related to the Experiments	 149
A.1 Information Note and Informed Consent Form	150
A.2 Questionnaire	152
A.3 JSON File	153
A.4 CSV File	155
A.5 Ethics Approvals	156
 B Résumé en français	 159
B.1 Introduction	159
B.2 Dispositif Expérimental	161
B.3 Impact de la Présence ou Absence de Compétition sur la Coopération Stig- mergique	162
B.3.1 Expériences	163
B.3.2 Comportements collectifs	164
B.3.3 Comportements individuels	164
B.3.4 Modèle de comportement des joueurs	167
B.4 Impact de la Composition du Groupe sur les Comportements Individuels et la Coopération Stigmergique	168
B.4.1 Expériences	168
B.4.2 Impact des <i>bots</i> sur la performance des participants	169
B.4.3 Modèle des stratégies de visites et de notations	169

Contents

B.4.4	Prédiction des profils comportementaux des participants humains . .	171
B.5	Impact de la Compétition Intragroupe et Intergroupe sur la Coopération	
	Stigmergique	172
	B.5.1 Expériences	172
	B.5.2 Performance et profils comportementaux des participants	174
	B.5.3 Modèle de prédiction des profils comportementaux individuels	174
B.6	Conclusion	177

List of Figures

1.1	Examples of stigmergy in animal societies	7
1.2	Examples of stigmergy in human groups	8
1.3	Dynamics of desire paths	9
2.1	Example of a grid with the corresponding cell values	19
2.2	Distribution of the cell values	20
2.3	Illustration of the cell selection and rating	21
2.4	Color scale	21
2.5	Evolution of the colors of the cells during the game	22
2.6	Player interface of the web application	25
2.7	Admin interface of the web application	26
2.8	Interface of the visualization application	28
2.9	Photographs of the experiments	29
2.10	Payment given to the participants	31
2.11	Participant demographics	32
3.1	Illustration of the bootstrap method	49
4.1	Probability to revisit a cell as a function of its value	54
4.2	Functions characterizing the rating probability distributions	56
5.1	Probability distribution functions of the normalized scores	62
5.2	Visit and rating performance	63
5.3	Human randomness bias	64
5.4	Inverse participation ratio	65
5.5	Fidelity	66
5.6	Values of visited cells	67
5.7	Probability of finding the cells with the highest values	68
5.8	Probability of revisiting cells visited in the previous round	69
5.9	Mean rating of the group	70
5.10	Rating probability	70
5.11	Individual mean rating on all games for Rule 1	72
5.12	Individual mean rating on all games for Rule 2	73
5.13	Behavioral profiles of individuals	74
5.14	Mean rating of each behavioral profile	75
5.15	Rating probability of each behavioral profile	76
5.16	Distribution of behavioral profiles as a function of the rank	77
5.17	Rating strategy of Mimic agents	80
5.18	Impact of the number of rounds on individual performance and collective dynamics	82

5.19 Impact of the group size on the collective performance and the dynamics of collective exploration and ratings	83
5.20 Impact of the rating strategy on agents' performance and the fidelity of rating	85
5.21 Rating probabilities of Optimized agents	87
5.22 Collective performance and dynamics of collective exploration and ratings for Opt-1 agents	88
5.23 Collective performance and dynamics of collective exploration and ratings for Opt-2 agents	89
5.24 Collective performance and dynamics of collective exploration and ratings for Opt-3 agents	90
5.25 Collective performance and dynamics of collective exploration and ratings for Opt-4 agents	91
5.26 Individual behaviors in experiments in which individuals play alone	93
5.27 Mean rating and behavioral profiles in solo games	94
6.1 Rating strategy of the bots	101
6.2 Mean normalized score	102
6.3 Individual performance and behavior of human participants in experiments with collaborator and defector bots	104
6.4 Individual performance and behavior of human participants in experiments with constant bots	105
6.5 Individual performance and behavior of human participants in experiments with optimized bots	106
6.6 Individual performance and behavior of human participants in experiments with groups of five humans	107
6.7 Probability of finding the cells with the highest values	108
6.8 Rank of the human participants	109
6.9 Behavioral profile of human participants	110
6.10 Rating probabilities of defectors in each experimental condition	112
6.11 Rating probabilities of Mimic agents	113
6.12 Two cues available to human participants	115
6.13 Performance of the PI and PIR models	118
6.14 Behavioral profiles given by the PI model	120
7.1 Performance of the participants	127
7.2 Distribution of collaborators	128
7.3 Performance of the Intra model and the Intra + Inter model	132
7.4 Performance of the Intra + Rank model	134
B.1 Exemple de tableau	161
B.2 Photographies des expériences	163
B.3 Probabilité de trouver les cellules ayant les valeurs les plus élevées	165
B.4 Note moyenne donnée par chaque profil comportemental	166
B.5 Profil comportemental des participants humains	170

List of Figures

B.6	Performance du modèle avec <i>(i)</i> et <i>(ii)</i> et du modèle avec <i>(i)</i> , <i>(ii)</i> et <i>(iii)</i> . .	172
B.7	Performance du modèle avec <i>(i)</i> et <i>(ii)</i>	176

List of Tables

3.1	Binning table	46
5.1	Percentages of each behavioral profile	76
5.2	Rating strategy parameters of Mimic agents	79
5.3	Visit strategy parameters of Mimic agents	79
5.4	Rating strategy parameters of Optimized agents	86
5.5	Visit strategy parameters of Optimized agents	86
6.1	Visit strategy parameters of the bots	101
6.2	Rating strategy parameters of the bots	101
6.3	Behavioral profile of human participants	110
6.4	Rating strategy parameters of Mimic agents	111
6.5	Visit strategy parameters of Mimic agents	113
6.6	Error between experimental and simulated observables	114
6.7	Parameters values of the model for predicting individual behavioral profiles	118
7.1	Distribution of behavioral profiles	129
7.2	Values of the potential features	132
A.1	CSV file example	155
B.1	Profils comportementaux des participants	166
B.2	Distribution des profils comportementaux	175

List of Definitions

3.1	Normalized group score	37
3.2	Mean normalized group score	38
3.3	Fractions of visits	38
3.4	Fractions of stars	38
3.5	Visit performances	39
3.6	Rating performances	39
3.7	Inverse Participation Ratio (IPR)	40
3.8	Fidelity of two vectors	41
3.9	Normalized individual score	42
3.10	Mean normalized score	43
3.11	Values of visited cells	43
3.12	Values of the cells visited with the highest values since the beginning	43
3.13	Probability of revisiting cells visited in the previous round	44
3.14	Probability of finding cells with given values	44
3.15	Mean rating	46
3.16	Intercept and slope parameters of the mean rating	47
3.17	Fractions of ratings	47

Remerciements

Je tiens à remercier sincèrement toutes les personnes qui ont contribué, de près ou de loin, à la réalisation de cette thèse.

Bien que ma passion pour la science remonte à mon enfance, un événement en 2018 m'a véritablement mené à écrire ce document que vous tenez entre vos mains (ou du moins que vous avez sous vos yeux).

Alors que je terminais ma première année au magistère de physique fondamentale d'Orsay, je suis tombé dans ma boîte aux lettres sur le numéro 537 du magazine La Recherche, intitulé *Chaos et systèmes complexes*. Ce magazine, auquel été abonnée l'une de mes colocataires, avait une couverture intrigante ornée d'un banc de poissons en mouvement. Cela a immédiatement capté mon attention. En plongeant dans son contenu, j'ai découvert le monde fascinant des systèmes complexes. Ce fut alors une révélation.

Cette découverte m'a conduit à m'orienter vers le master international *Physics of Complex Systems*, puis à effectuer un stage sur les réseaux complexes avec Alain Barrat et Márton Karsai. Finalement, cette thèse, menée sous la direction de Clément Sire et Guy Theraulaz, en est la suite logique.

La plupart des chercheurs ayant contribué à ce dossier du magazine La Recherche ont joué un rôle important tout au long de mon parcours. Parmi eux, on retrouve : Clément Sire et Guy Theraulaz, mes deux directeurs de thèse ; David Chavalarias, l'un des rapporteurs de cette thèse ; Chiara Poletto, qui m'avait accepté pour une thèse sur la modélisation de la propagation des épidémies seulement deux mois avant la pandémie de Covid-19 ; Florent Krzakala, l'un de mes professeurs de master qui m'avait également proposé une thèse ; Lenka Zdeborová, Pablo Jensen et Marc Barthelemy, trois chercheurs que j'avais contactés pour des opportunités de doctorat ; et enfin Christian Jost, un chercheur de mon équipe au CRCA. Leur travail et leur influence ont tous, d'une manière ou d'une autre, façonné mon parcours.

Ainsi, je tiens à remercier Ninon, mon ancienne colocataire qui ne lira jamais ces lignes, sans qui je n'aurais probablement pas découvert ce domaine de la physique. D'ailleurs, étant donné que j'ai encore le magazine, je me demande si j'en ai acheté un autre ou si je ne te l'ai jamais rendu. Quoi qu'il en soit, merci de m'avoir indirectement mis sur cette voie.

Venons-en aux remerciements pour ceux qui ont directement travaillé avec moi durant cette thèse.

Tout d'abord, merci à Maxime Delacroix, notre grand stagiaire, pour avoir effectué toutes les modifications sur l'application web STIGMER. Tu as su t'approprier un sujet dont tu ne connaissais rien et améliorer une application mal documentée et pas toujours bien codée. Merci également pour ton implication qui a largement dépassé le cadre de ton stage initial. J'ai toujours pu compter sur toi pour m'aider lorsque l'application ne fonctionnait plus sans que je ne sache pourquoi. Ce fut un plaisir de travailler avec toi

et de pouvoir contribuer à la fin de ta formation en te transmettant mes connaissances en programmation. Au plaisir de partager avec toi de nouvelles séances d'observations du ciel.

Merci Stéphane Cezera pour l'organisation des expériences et le recrutement des participants. Pendant les expériences, c'est toi qui avais la mission la plus dure : trouver, par temps froid et pluvieux, des remplaçants pour les participants qui manquaient à l'appel. Merci Adrien Blanchet pour tes contributions et ton aide dans les expériences. J'aurais aimé travailler davantage avec toi sur des sujets plus théoriques, mais le temps d'une thèse est limité. Eskerrik asko Ramón Escobedo pour tes contributions lors de nos réunions hebdomadaires et tes précieux conseils (notamment pour supporter les humeurs de Clément et Guy).

Je tiens également à dire un très grand merci à Clément Sire et Guy Theraulaz, mes deux directeurs de thèse, pour leur accompagnement tout au long de ces trois années et demie. Vous m'avez offert la possibilité de travailler sur un sujet très intéressant, challengeant et impliquant tous les aspects de la recherche : du design expérimental à l'expérimentation, en passant par l'analyse de données, la modélisation et la simulation. Vous m'avez appris que l'interdisciplinarité n'est pas toujours évidente. Il n'est pas facile de faire collaborer des personnes de perspectives différentes, mais c'est précisément cette diversité qui a enrichi ma thèse et mon parcours. Merci également pour votre aide précieuse dans les dernières semaines de rédaction, rendant ce document bien plus complet et lisible. Je retiendrai surtout vos deux tempéraments distincts et forts qui ont rendu nos réunions si dynamiques et mouvementées pour ne pas en dire plus. J'espère que vous avez apprécié travaillé sur cette thèse autant que moi. J'espère également que ce projet continuera d'une manière ou d'une autre.

Je souhaite aussi remercier David Chavalarias et Laëticia Gauvin, les rapporteurs de cette thèse, ainsi qu'Alain Barrat et Umberto Grandi, les deux autres membres du jury. Merci pour vos retours constructifs sur mon travail de recherche, qui ont grandement contribué à son amélioration. Votre déplacement jusqu'à Toulouse pour assister à la soutenance a été un honneur et un soutien précieux pour moi.

Enfin, je tiens à remercier le CNRS pour avoir entièrement financé ma thèse à travers plusieurs moyens : l'appel à projet 80|Prime soutenant des projets de recherche interdisciplinaires multi-équipes (3 ans), l'institut de physique (2 mois), et la délégation Occitanie Ouest (2 mois). Leur soutien a été crucial pour la réalisation de mes travaux.

Il est temps de rendre hommage aux nombreux chercheurs et doctorants que j'ai eu le privilège de côtoyer tout au long de cette aventure.

Je souhaite débiter en exprimant ma gratitude à l'ensemble des membres du LPT que j'ai côtoyés quotidiennement. En particulier, je voudrais remercier chaleureusement Nicolas Destainville, un très bon directeur de laboratoire, qui a également été mon maître de stage lors de ma première année de master. Un immense merci à Malika Bentour, qui a géré tous les aspects administratifs liés à mes contrats, voyages, et à l'organisation des expériences au LPT. Son travail a grandement facilité mon quotidien. Enfin, je tiens à remercier Sandrine Le Magoarou, qui a consacré un temps non-négligeable à nos côtés pour le déploiement de l'application web STIGMER sur les serveurs du LPT, veillant à ce

List of Definitions

que tout fonctionne parfaitement. Je tiens aussi à remercier les membres du CRCA, qui même si je ne les ai pas beaucoup côtoyés, ont toujours été d'une grande gentillesse.

Je ne saurais oublier les nombreux doctorants, post-doctorants et stagiaires qui ont défilé au cours des années et avec qui j'ai partagé la majorité de mes journées, ainsi que d'innombrables repas et pauses café. Je ne pourrai pas nommer tout le monde, mais je remercie tout particulièrement ceux dont j'ai été le plus proche : Jordan Heraud, Julie Cornet, Faedi Loulidi, Théo Hennequin, Claire Lacouture, Naïmo Davier, Denis Rochette, Thomas Reppinger, Noam Izem et Quentin Bedel. Je n'oublierai jamais nos nombreuses discussions, soirées et surtout notre super week-end ski. J'espère avoir l'occasion de vous revoir tous bientôt. Bien entendu, je remercie également tous les autres que je n'ai pas cités, mais avec qui j'ai partagé d'agréables moments au cours de ma thèse.

Je tiens ensuite à exprimer ma profonde gratitude envers NanoX pour leur généreuse bourse qui m'a permis de participer pendant trois semaines à la *Complex Systems Summer School* à Santa Fe, au Nouveau-Mexique, aux États-Unis. Cette école d'été a indéniablement été l'un des moments forts de ma thèse. J'ai eu l'opportunité d'y rencontrer de nombreux doctorants, post-doctorants, chercheurs, ainsi que des journalistes scientifiques et entrepreneurs évoluant dans le domaine des systèmes complexes. Cette expérience m'a offert l'occasion d'explorer de nouvelles applications de ces systèmes complexes, tout en contribuant à plusieurs projets scientifiques passionnants que j'aurais aimé approfondir davantage. Je tiens donc à remercier chaleureusement tous les participants, organisateurs et conférenciers pour avoir rendu cette expérience intéressante, enrichissante et mémorable.

Pour finir, je tiens à remercier toutes les personnes extérieures au domaine de la recherche qui ont contribué à rendre ces années si agréables à vivre.

Je souhaite commencer par mes parents, qui ont toujours été là pour moi et sans qui je n'aurais pas été là pour écrire cette thèse. Merci pour votre soutien inébranlable et pour avoir toujours cru en moi. Je tiens également à remercier mon frère, Sacha, qui a été mon colocataire pendant plus d'un an et demi au début de la thèse. J'espère que ta nouvelle aventure chez Ankama se passe bien, et que tu pourras t'y épanouir pour un moment. Il me tarde de jouer aux jeux dont tu auras participé à la réalisation. Ton talent et ta passion me rendent fier d'être ton frère.

Je tiens tout particulièrement à remercier Camille pour m'avoir supporté et accompagné tout au long de cette thèse. Ton amour et ta patience, surtout durant les derniers mois de rédaction, ont été essentiels pour moi. Je sais que ces périodes n'ont pas été faciles, mais ton soutien m'a donné la force de continuer. Merci pour le temps que tu as passé à la relecture de ma thèse et pour les quelques schémas que tu as dessinés. Merci de m'avoir écouté répéter ma soutenance et de m'avoir aidé à l'améliorer. Merci de me mettre un couteau sous la gorge en me dictant ces quelques mots. Je t'aime profondément et je suis reconnaissant pour tout ce que tu as fait pour moi.

Je remercie également Malo. Bien que ta présence ait parfois entraîné quelques désagréments, comme le fait que je doive me moucher dix fois par jour et prendre des médicaments tout le temps, merci pour ton ronronnement qui a accompagné mes longues journées de télétravail et mes soirées. Merci pour tes papouilles et tes pipouilles.

Enfin, je voudrais remercier tous mes amis et en particulier ceux du club de Portet

Gym. Faire de la gym et boire des bières à vos côtés m'ont permis de me défouler et de passer du bon temps tout au long de ces quinze dernières années.

Finalement, je voudrais exprimer ma gratitude envers toutes les personnes qui, par le biais de leurs œuvres, m'ont accompagné pendant ces quelques années. Que ce soit à travers des livres inspirants (Ted Chiang, Alain Damasio, Frank Herbert, Liu Cixin, Stephen King...), des bandes dessinées et mangas fascinants (Naoki Urasawa, Mathieu Babellet...), des albums de musique captivants (Animals as Leaders, GoGo Penguin, Immortal Onion, Sofiane Pamart, Thom Yorke...), des concerts mémorables (Dream Theater, French 79, GoGo Penguin, Jungle, Moon Hooch, Nils Frahm, Nothing But Thieves, Rone, The Smile...), des podcasts enrichissants (*La Méthode scientifique* devenu *La Science*, *CQFD*, *C'est plus que de la SF*, *FloodCast*...), ou des vidéos postés sur YouTube (3Blue1Brown, Numberphile, Veritasium, Vsauce...), vous avez tous contribué à nourrir ma curiosité, à me motiver, et à m'évader le temps de quelques instants. Merci pour votre talent et votre passion.

Je tiens également à remercier toutes les autres personnes qui ont fait le déplacement pour assister à cette soutenance, témoignant de leur intérêt et de leur soutien.

Merci à ceux que j'ai oubliés.

Merci à tous.

Introduction

Chapter Summary

Collective intelligence refers to the capacity of communities to collaboratively address challenges, make decisions, and generate insights that surpass individual capabilities. Collective intelligence is based on two phenomena: social information and social influence. Social information is the knowledge and cues that individuals in a group share with each other. Social influence is the process by which individuals adjust their attitudes, behaviors, or opinions in response to social situations.

The concept of collective intelligence finds practical application through mechanisms such as stigmergy which is commonly found in animal and human societies. Stigmergy is a coordination mechanism in which traces left by individuals in a medium guide and stimulate their subsequent actions. Examples of stigmergy include ant foraging behavior, wasp comb building, and desire path creation. However, the digitization of human society has given rise to a new form of stigmergy that relies on digital traces. Examples of such stigmergic processes notably include rating and recommender systems that extensively use the digital traces left by their users.

Cooperation is the process by which groups of organisms work together for common or mutual benefits. The prevalence of cooperation among social species raises questions about its evolution in a world where natural selection can favor selfishness. Several mechanisms have been proposed to explain this evolution, including kin selection, reciprocity, network reciprocity, and group selection. Stigmergy also provides insights into unintended cooperation through indirect interactions. Competition, which operates at various levels, is another critical factor in the development of cooperation. The rivalry between individuals and groups can be a catalyst for improvement, but it can also create incentives for unethical or deceptive behavior.

Numerous strategies have been proposed and tested to enhance collective intelligence in human groups. One way to do this is through changes in group size and composition, or through changes in the network structure of interactions. Another option is to improve social information exchange by increasing control over it and reducing noise. Social nudges and bots can also be utilized to influence behaviors.

In the digital age, social information in the form of digital traces has a significant impact on human decision-making. It is crucial to gain a deeper comprehension of the societal impact of these traces, especially considering their vulnerability to manipulation. In this context, this Ph.D. thesis aims to investigate how humans use and interact with digital traces, with a focus on the conditions for cooperation and the influence of competition on their utilization.

Contents

1.1	Introduction	2
1.2	Collective Intelligence in Human Groups	3
1.2.1	Definition	3
1.2.2	Types of collective problem-solving	3
1.2.3	Social information	4
1.2.4	Sharing social information	4
1.2.5	Social influence	5
1.3	Stigmergic Problem-Solving	5
1.3.1	Stigmergy as a universal coordination mechanism	5
1.3.2	Different types of stigmergy	6
1.3.3	Stigmergic processes in bacterial and animal societies	6
1.3.4	Stigmergic processes in human groups	8
1.4	Cooperation and Competition	10
1.4.1	Cooperation in human groups	10
1.4.2	The evolution of cooperation	10
1.4.3	The dual nature of competition	11
1.5	Enhancing Collective Intelligence	12
1.5.1	Exchange of social information	12
1.5.2	Size and composition of the group	13
1.5.3	Network structure	14
1.5.4	The role of noise in group decision-making	14
1.5.5	Social nudges and bots	15
1.6	Outline of the Thesis	15

1.1 Introduction

In today's interconnected world, human interactions create complex patterns [Helbing, 2012; Ball, 2012; Gauvin et al., 2009; Barrat et al., 2008]. Like ants navigating their nest or birds gracefully flocking across the sky, human groups exhibit a remarkable capacity for coordination and cooperation. These collective behaviors, often characterized by their emergent properties, have captured the imagination of researchers and thinkers for centuries [Camazine et al., 2001; Sumpter, 2010; Dugatkin, 2002]. As we move deeper into the digital age, our capacity to observe, analyze, and influence these behaviors has reached unprecedented levels, offering new insights into the foundations of collective intelligence in human societies [Stonier, 1992].

At the heart of this exploration is the concept of complex systems, which is relevant to fields as diverse as biology, sociology, physics, and economics [Ball, 2006]. Complex

1.2. Collective Intelligence in Human Groups

systems are systems in which a large number of individual components interact, often in nonlinear ways, giving rise to behaviors that cannot be easily predicted from the properties of individual elements alone. The emergent properties of these systems, where the whole is greater than the sum of its parts, provide a powerful means of understanding the complexity of human groups.

Emergence refers to the enigmatic phenomenon by which new patterns, structures, or behaviors emerge from the interactions of individual agents within a system [Nicolis and Nicolis, 2012]. Consider, for example, the murmuration, in which thousands of birds display synchronized movements across the sky, seemingly orchestrated by an imperceptible force [King and Sumpter, 2012]. Notably, no single bird possesses the knowledge or orchestrates this remarkable show. Yet, together, they produce a mesmerizing manifestation of collective intelligence. This emergence is not limited to the natural world; it extends to the domain of human societies and the behaviors that emerge from the interactions of people within groups [Ball, 2012; Goldstone et al., 2008].

Collective behaviors are the manifestations of these emergent properties in societies. They cover a range of activities from problem-solving and decision-making to information aggregation, and opinion formation. I will use this framework of collective behaviors to unravel the complexity of collective intelligence in human societies.

1.2 Collective Intelligence in Human Groups

1.2.1 Definition

Collective intelligence represents the ability of communities to collaboratively tackle challenges, reach decisions, and generate insights surpassing individual capabilities [Bonabeau, 2009; Salminen, 2012; Woolley et al., 2015; Malone and Bernstein, 2015; Gunasekaran et al., 2016]. Collective intelligence is characterized by the collaborative efforts of individuals within a group that lead to the emergence of shared knowledge, innovative solutions, and effective decisions. It goes beyond the mere aggregation of individual knowledge and involves a complex combination of different points of view and information processing mechanisms in social settings.

1.2.2 Types of collective problem-solving

Collective intelligence encompasses a variety of problem-solving approaches designed to address specific challenges. This non-exhaustive section explores several categories of collective problem-solving.

One category is estimation tasks, where individuals provide estimates that reveal how diverse judgments aggregate into a collective estimate [Jayles et al., 2017]. Another is opinion and knowledge aggregation, where a group must gather, aggregate, and synthesize opinions or knowledge to reach a unified perspective [Kameda et al., 2022]. Idea collection emphasizes collaborative idea generation, fostering creativity to solve problems [Bothos et al., 2009, 2012; Goncalo and Duguid, 2012]. Forecasting benefits from collective intelligence, providing insight into forecast accuracy and reliability [Sjöberg, 2009; Ball

and Ghysels, 2018]. Consensus decision-making is the process in which individuals work through conflicting perspectives to reach a unanimous decision [Dyer et al., 2008; Kameda et al., 2012; Santos et al., 2016; Bose et al., 2017; Marshall et al., 2019; Kameda et al., 2022]. Finally, coordination optimizes productivity between tasks, whether independent or interdependent, to increase group efficiency [Kittur et al., 2009].

This comprehensive exploration illustrates the versatility of collective intelligence in different contexts and highlights its ability to adapt to a wide range of challenges.

1.2.3 Social information

Social information, defined as the shared knowledge and cues exchanged among individuals within a group, serves as a cornerstone in the coordination of group activities and collective problem-solving [Camazine et al., 2001; Garnier et al., 2007; Couzin, 2007; Goldstone and Gureckis, 2009; Moussaid et al., 2009]. It is the means by which a collective functions cohesively and efficiently. The importance of social information lies in its ability to accelerate adaptive responses beyond the effectiveness achievable through individual information gathering alone [Zentall and Galef Jr., 1988; Heyes and Galef, 1996; Danchin et al., 2004; Laland, 2004; Duboscq et al., 2016].

The evolutionary advantages conferred by the utilization of social information are particularly pronounced in the animal kingdom, where coordinated efforts are crucial for survival and reproduction. This extends to diverse contexts, including foraging, decision-making, division of labor, nest building, and colony defense [Camazine et al., 2001; Garnier et al., 2007; Sumpter, 2010; Seeley, 2010]. The exchange of social information in these settings is a strategic adaptation that enhances the collective competence of social groups.

1.2.4 Sharing social information

The dissemination of social information within a group involves several modes, each of which influences the dynamics of collective intelligence.

Direct information sharing occurs through explicit communication channels that allow individuals to transparently communicate knowledge, preferences, or insights to others. This direct exchange is characterized by the clarity and immediacy of the information conveyed, fostering a shared understanding among group members.

In contrast, indirect sharing of social information involves subtle cues, observational learning, or implicit signals that shape individual behavior and decision-making processes within the group [Baltzersen, 2022]. Individuals may adjust their actions based on observed behaviors or outcomes of others, resulting in a form of implicit knowledge transfer. This indirect sharing mechanism is particularly prominent in situations where explicit communication is limited or in contexts where nonverbal cues play a critical role.

In most situations, however, social information exchange is a mixture of direct and indirect exchange. These hybrid modes of information transmission often involve a combination of explicit communication and observational learning. They contribute to the development of a robust information network that enhances the adaptive capabilities of the collective.

1.3. Stigmergic Problem-Solving

The use of social information varies from person to person. While some people incorporate it into their decision-making process, others may entirely ignore it and rely solely on their private information. There are even cases where individuals intentionally resist the message conveyed by social information [Jayles et al., 2017]. Moreover, the same person may adapt their approach to providing and using social information based on the context at hand, reflecting the dynamic nature of individual decision-making and information processing [Steinel et al., 2010].

1.2.5 Social influence

Social influence is a fundamental aspect of group dynamics, exerting a profound impact on individual decision-making processes and shaping the emergent properties observed at a group level [Asch, 1955; Cialdini and Goldstein, 2004; Cialdini, 2007; Moscovici, 2014]. It encompasses the ways in which individuals adapt their attitudes, behaviors, or opinions in a social context.

Social influence manifests itself in various ways and includes phenomena such as conformity, peer pressure, compliance, and persuasion. Conformity is the tendency of individuals to adjust their attitudes, behaviors, or beliefs to reflect those of a majority or societal norm, even in the absence of explicit pressure. Peer pressure is similar to conformity, but is specifically exerted by individuals who are similar in age, gender, or shared experiences. Compliance, on the other hand, involves individuals adjusting their behavior in response to a direct request or command, often driven by a desire for approval or to avoid confrontation. Finally, persuasion involves influencing others through the use of rhetoric and communication strategies.

Numerous factors play a role in the dynamics of social influence, including the strength of interpersonal relationships, the credibility of information sources, and the susceptibility of individuals to conform to or deviate from group norms. Importantly, social influence is not inherently negative, and its effects depend on the specific context and underlying motivations at play.

1.3 Stigmergic Problem-Solving

1.3.1 Stigmergy as a universal coordination mechanism

While studying the coordination mechanisms of termites, Pierre-Paul Grassé introduced the concept of stigmergy in 1959 [Grassé, 1959; Baltzersen, 2022]. Stigmergy is based on the fundamental principle that traces left by an individual in the environment serve as informational cues that elicit responses from other individuals, whether the same or different [Theraulaz and Bonabeau, 1999; Giuggioli et al., 2013; Heylighen, 2016a]. The term “stigmergy” is etymologically rooted in the Greek words “stigma”, meaning “sign”, and “ergon”, meaning “work”, and more generally, “action”, indicating the intrinsic link between the signs left by individuals and the subsequent actions of others.

In stigmergic processes, individuals appear to be doing things on their own, almost randomly. However, when observed as a collective, coherent and coordinated behavior

emerges. Stigmergy is deeply rooted in the notion of indirect interactions between individuals within a group, leading to the spontaneous emergence of collective actions free from centralized control or open communication. Within this framework, traces function as an external memory that allows individuals to inform their decisions and actions [Theraulaz and Bonabeau, 1995; Garnier et al., 2007; Khuong et al., 2016; Collignon and Detrain, 2021].

1.3.2 Different types of stigmergy

Stigmergy can be divided into four distinct categories based on two key distinctions: quantitative vs. qualitative and marker-based vs. sematectonic [Heylighen, 2016b; Baltzersen, 2022; Van Dyke Parunak, 2006].

Quantitative stigmergy involves actions whose intensity and frequency depend on the strength of the trace left by the individual. In simple terms, a stronger trace results in more intense actions, similar to the concept of a potential field in physics. Qualitative stigmergy works differently. Actions are no longer determined by the intensity of the trace, but by specific trace characteristics or environmental conditions that trigger different actions [Theraulaz and Bonabeau, 1999].

Another important distinction lies in the nature of the trace itself. Marker-based stigmergy occurs when individuals intentionally leave specific markers as signals for others. Conversely, sematectonic stigmergy occurs when there are no explicit markers; instead, the current state of the environment itself serves as a trace that influences subsequent actions.

1.3.3 Stigmergic processes in bacterial and animal societies

Manifestations of stigmergy within animal societies are most prominent in social insects, including ants, termites, bees, and other [Theraulaz and Bonabeau, 1999]. Among these, the foraging behavior of ants stands out as an iconic illustration of stigmergic principles in action (see Figure 1.1A). When ants independently discover different food sources, stigmergic mechanisms come into play, as these insects effectively select the closest and most accessible source to the nest [Goss et al., 1989; Deneubourg et al., 1990; Beckers et al., 1990]. The pheromone left by the ants is more important when the food source is closer, so more ants will go to that source, amplifying the pheromone. This complex process ultimately selects the optimal path.

The sequential comb building of social wasps is another great example of stigmergy, where the configuration of the comb itself provides enough information to ensure the coordination of the wasps' building activity (see Figure 1.1B) [Theraulaz and Bonabeau, 1995, 1999; Khuong et al., 2016]. Indeed, as the wasps add cells to the comb, the choice of cell location is not arbitrary but is intricately linked to the current configuration of the comb. The more walls already built, the more likely a wasp is to choose that particular location to build a cell. This preference for sites with more walls introduces a qualitative dimension to the process

Stigmergy is also present in solitary species, as exemplified by the nest construction

1.3. Stigmergic Problem-Solving

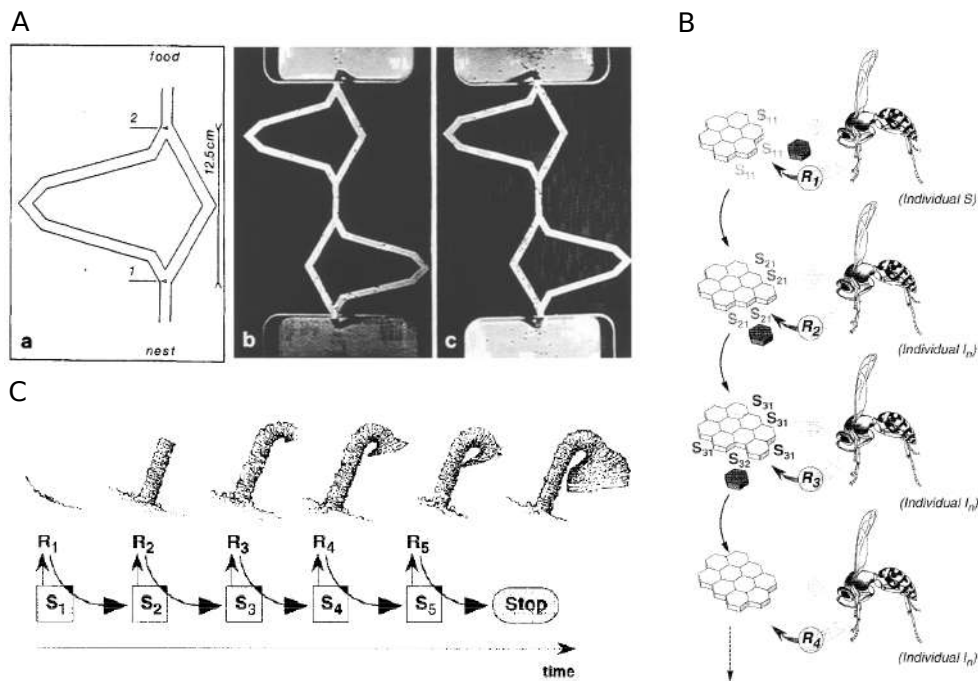


Figure 1.1: **Examples of stigmergy in animal societies.** (A) Foraging behavior of ants. At the beginning of the experiment (picture b), the ants explore all possible paths to search for food. However, over time (picture c), they converge onto the shortest path. (B) Comb building of social wasps. Each wasp accesses the current configuration of the comb and adds a new cell to the location with the most walls already constructed. By doing so, the wasps alter the configuration of the comb, which affects the location of the following cells. (C) Nest construction of the solitary wasp *Paralastor* sp. The wasp constructs its nest in several stages, each triggering the next one. Credits: (A) [Goss et al., 1989], (B) and (C) [Theraulaz and Bonabeau, 1999].

of the wasp *Paralastor* sp (see Figure 1.1C) [Theraulaz and Bonabeau, 1999]. The nest construction process unfolds in discrete stages, with each stage triggering a specific action by the wasp. Remarkably, the wasp does not require a premeditated plan or memory for this task, as it can determine its current stage by examining the work it has already completed. This lack of memory becomes apparent when the sequence of stages is disrupted. For example, if an observer makes a hole that triggers the first stage on an almost completed nest, the wasp will not finish the nest but will reset its construction process and start again from the beginning [Smith, 1978]. This intriguing experiment highlights the concept of stigmergy, where the wasp’s actions are closely linked to the current state of the environment.

While stigmergy has received attention primarily in the context of social insects, stigmergy is also prevalent in bacterial communities [Gloag et al., 2015]. Bacteria use chemical cues (chemotaxis), such as quorum sensing molecules, to modulate their behavior in response to collective actions [Li and Tian, 2012]. Bacteria also exhibit trail-following behavior through the use of sematectonic and quantitative stigmergy [Gloag et al., 2013].



Figure 1.2: **Examples of stigmergy in human groups.** (A) Example of desire path showcasing an alternative route where individuals naturally traverse the roundabout instead of following the designated pathway. (B) Revision history that describes the additions to and deletions from a Wikipedia article. (C) Methods used by three websites (Amazon, TripAdvisor, and IMDb) to display user ratings. Credits: (A) <https://imgur.com/JIXACph>.

1.3.4 Stigmergic processes in human groups

Stigmergy, which originated in the study of social insects, has found increasingly relevant applications in understanding human behavior [Doyle and Marsh, 2013; Baltzersen, 2022; Van Dyke Parunak, 2006]. While some forms of social information exchange in human groups are direct, such as face-to-face conversations, human communication is mostly indirect, much like the traces ants leave in their environment. Centuries of human history attest to the profound impact of these traces on society, from handwritten notes and copied books to the modern digital age with its disruptive technologies.

In non-digital settings, a prime example of sematectonic quantitative stigmergy is the phenomenon of “desire paths” in urban environments, where people take shortcuts across grassy areas instead of following designated paths (see Figure 1.2A). These desired paths, etched by individuals seeking more efficient routes, guide the behavior of others, gradually reinforcing and formalizing the paths as stigmergic markers [Helbing et al., 1997]. City planners often recognize the natural development of these informal pathways and intentionally incorporate them into formal infrastructure to align with the organic flow of human movement (see Figure 1.3).

The proliferation of digital technologies presents new challenges and opportunities for human adaptation [Cochoy et al., 2017; Golder and Macy, 2014]. In online environments, human behavior has been transformed into a form of digital stigmergy [Doyle and Marsh, 2013; Baltzersen, 2022; Van Dyke Parunak, 2006]. Individuals now leave digital traces, often referred to as data, across the vast expanse of the Internet. These traces encompass a spectrum of activities, from passive records such as web cookies and browsing history to active engagements such as social media posts, comments, likes, and shares. These digital traces serve as markers that guide subsequent actions, shape user behavior, and contribute

1.3. Stigmergic Problem-Solving



Figure 1.3: **Dynamics of desire paths.** Satellite images of the entrance area of the Université de Toulouse III – Paul Sabatier. (A) 2006: Absence of any discernible desire path. (B) 2013: A desire path appears after the construction of a metro and bus hub in the upper left corner. (C) 2023: The university builds an official path following the desire path. However, the introduction of Téléo, a cable car line in the lower left corner, leads to the formation of new desire paths. Credits: Google Earth.

to the emergence of collective responses.

Wikipedia, the collaborative online encyclopedia, is an example of digital stigmergy in action [Loveland and Reagle, 2013; Rezugui and Crowston, 2018]. Contributors from around the world create and edit articles, leaving traces of their work in the form of revisions and discussions (see Figure 1.2B). These traces guide subsequent edits and shape the evolution of articles, illustrating how digital markers influence collective knowledge-building efforts. This is an example of sematectonic qualitative stigmergy.

Online platforms such as TripAdvisor, eBay, Amazon, and Uber provide a rich landscape for exploring marker-based quantitative stigmergy. These platforms rely on recommendation and rating systems to help users make informed decisions. Users actively participate by providing ratings, reviews, and feedback, each of which acts as a marker (see Figure 1.2C). High ratings and positive feedback influence others to choose certain products, services, or experiences. Understanding this aspect of digital stigmergy, and in particular how digital markers in the form of ratings and reviews guide individual decisions and foster a collective response that shapes the dynamics of these online ecosystems, is the focus of this Ph.D. thesis [Lü et al., 2012; Hennig-Thurau et al., 2012; Jesse and Jannach, 2021].

However, the use of digital traces is very sensitive to noise and manipulation [Herlocker et al., 2004; Gunes et al., 2014; Grandi and Turrini, 2016]. 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.

1.4 Cooperation and Competition

1.4.1 Cooperation in human groups

Cooperation is the process where groups of organisms work together for common or mutual benefits [Axelrod, 2006; Rand and Nowak, 2013]. It is a ubiquitous phenomenon that can be observed among individuals of the same species [Dugatkin, 1997], and even between individuals of different species, for example, via symbiosis or mutualism [Bronstein, 2015]. However, the highest levels of cooperation are rare and are only found in social species, such as social insects or human societies. But cooperation in human groups is even more surprising because of its scale, both in terms of the number of individuals involved and the degree of cooperation [Melis and Semmann, 2010]. These altruistic acts, occasionally extending to strangers who will never see each other again, may sometimes defy intuition, especially considering that genuine altruism often entails a personal cost to the altruistic individual [Kitcher, 1993; Fehr and Fischbacher, 2003; Feigin et al., 2014; Wilson, 2015].

Cooperation, and more precisely its evolution, is therefore considered quite paradoxical. How such a phenomenon could have evolved in a world in which natural selection introduces competition which often favors selfishness? Hence, there must be other processes at play that favor the evolution of cooperation above the simpler selfish acts.

1.4.2 The evolution of cooperation

The question of the evolution of cooperation has long been the focus of many studies [Axelrod, 2006]. In 2006, Novak summarized these mechanisms into five categories, shedding light on the various ways in which cooperation can emerge and persist in evolutionary contexts [Nowak, 2006].

- *Kin selection* is the idea that individuals may favor cooperation when it benefits their genetically related relatives. This concept is based on the notion that genes that promote cooperative behavior can spread indirectly through shared family ties, ultimately leading to the evolutionary success of cooperative traits. This mechanism is often observed in social organisms such as bees or ants, where individuals work together to support their genetic line.
- *Direct reciprocity* is well summarized by the proverbial expression “You scratch my back, and I’ll scratch yours”. In this scenario, individuals cooperate with others on a tit-for-tat basis, providing help or resources to those who have previously helped them. This creates an ongoing cycle of cooperation and mutual support within social groups. Direct reciprocity is particularly prominent in species with long periods of interaction and the ability to remember past interactions, such as certain primates and humans.
- *Indirect reciprocity* introduces a more complex dimension to cooperation, often referred to as “reputation-based”. In this context, individuals cooperate not only based on personal interactions but also consider the reputation of others. They tend to favor cooperation with individuals who have a history of cooperative behavior, and

1.4. Cooperation and Competition

this system relies on the transmission of information about the behavior of individuals within a group. This mechanism is present in human societies and increasingly recognized as crucial for shaping cooperation in large, anonymous populations [Panchanathan and Boyd, 2004; Santos et al., 2021].

- *Network reciprocity* extends the idea of cooperation within a particular group or network. In this framework, cooperation is more likely to be sustained when it occurs among individuals who share common interactions and connections. Such networks create clusters of cooperators who benefit from working together. Network reciprocity is evident in various contexts, such as social alliances in non-human species and cooperative behavior in online communities or social groups.
- *Group selection* introduces a broader perspective, emphasizing that cooperative groups can outperform non-cooperative ones at the group level [Wilson, 2015]. In this scenario, groups with a higher prevalence of cooperation have a better chance of survival and reproduction, contributing to the long-term persistence of cooperative traits. Although group selection has been controversial in evolutionary biology, it remains an important mechanism for understanding how cooperation evolves in different species and societies [Richerson et al., 2016].

These mechanisms that promote cooperation under individual-level competition interact with the competition that occurs at other scales of biological organization. They can impact both the individual advantage of defection and the collective benefit of cooperation. Evolutionary game theory has been extended to describe natural selection operating at multiple levels of organization in a group-structured population. This includes competition among individuals within a group and competition among groups of individuals [Wilson, 1975; Wilson and Dugatkin, 1997; Traulsen and Nowak, 2006; Luo, 2014; Simon et al., 2013; Cooney, 2019]. This generalization provides a framework for exploring the countervailing effects of the individual incentive to defect and the collective incentive to cooperate, highlighting the evolutionary tug-of-war between levels of selection. In this framework, altruism evolves when between-group selection prevails over within-group selection [Wilson, 2015].

Over the years, stigmergy, described in the previous section, has also been proposed to explain cooperation. In stigmergy, indirect interactions can lead to unintended cooperation with individuals [Chiong and Kirley, 2012].

1.4.3 The dual nature of competition

As a fundamental aspect of human and animal behavior, competition plays a multifaceted role in shaping societies, driving innovation, and influencing the dynamics of cooperation [Society and Press, 1993; Axelrod, 1997; Beersma et al., 2003; Toma and Butera, 2015; Wilson, 2015]. In essence, competition can be defined as rivalry among individuals or groups for limited resources, opportunities, or advantages. It occurs when individuals or entities strive to outperform others to secure benefits that can range from tangible rewards to intangible recognition. A compelling example of competition in nature is

foraging, the search for and gathering of food. These competitive interactions often occur at multiple levels: within groups, between groups, and even between species, creating multilevel complexity.

Competition is often a catalyst for innovation and performance improvement [Gilbert, 2006]. The pursuit of an edge over rivals forces individuals and organizations to push their limits, explore new ideas, and adopt new strategies. This drive to outperform the competition can lead to breakthroughs that would not have been possible in the absence of competition. In the business sector, companies competing for market share invest in research and development, resulting in the creation of new products and services. This sustained cycle of innovation benefits consumers by offering them greater choices and higher-quality products.

However, when success is measured solely in terms of outperforming others, there is less inclination to share critical knowledge, resources, or support. Individuals may choose to share less information or even resort to sharing misinformation [Mokkonen and Lindstedt, 2016]. In addition, the competitive environment can create incentives for unethical or deceptive behavior [Kurvers et al., 2021]. This not only compromises the integrity of the competition but also undermines trust among participants, threatening the structure of the collaborative effort and lowering overall group performance and payoff [Barker et al., 2012].

1.5 Enhancing Collective Intelligence

Many strategies for enhancing collective intelligence in human groups have been proposed and tested. This section provides a non-exhaustive review of some of these strategies that influence group performance and collective intelligence, based on empirical research.

1.5.1 Exchange of social information

The exchange of social information, a fundamental aspect of collective intelligence, plays a critical role in shaping group decisions [Moussaid et al., 2009]. By fostering communication and information flow among individuals, this process leverages the diverse knowledge and perspectives within a group and can enhance collective intelligence and group performance.

The exchange of social information facilitates the incorporation of diverse perspectives within the group. Studies suggest that diverse groups exhibit enhanced problem-solving capabilities due to the variety of insights and approaches brought to the table [Hong and Page, 2004]. Moreover, studies have shown that social information can significantly improve estimation accuracy, highlighting its positive impact on decision-making processes [Jayles et al., 2017].

Furthermore, the exchange of social information fosters a collaborative environment where individuals can leverage the collective knowledge of the group. This collaborative synergy is exemplified in domains ranging from scientific research to business innovation, where information sharing sparks creativity and accelerates problem-solving processes [Goncalo and Duguid, 2012]. In essence, social information sharing acts as a catalyst

1.5. Enhancing Collective Intelligence

for the emergence of novel ideas and solutions that may not have been possible within the confines of individual knowledge.

Social information can therefore enhance collective intelligence, but one should consider the format of this information. Research suggests that presenting aggregated information tends to be more effective than providing individual pieces [Jayles et al., 2021b; Surowiecki, 2005; Mavrodiev et al., 2013]. In addition, studies suggest that due to biases inherent to human cognition, there are cases where presenting incorrect information as a strategic countermeasure against the effects of these biases may be more beneficial than providing correct information [Jayles et al., 2020b, 2021a; Kao et al., 2018].

However, social information sharing is not without its challenges. The risk of conformity, as demonstrated in Asch [1955], is a notable concern. Individuals may succumb to the opinions of the majority, potentially leading to suboptimal decisions. This conformity risk can be exacerbated by information cascades, famously illustrated by Milgram’s experiment [Milgram et al., 1969], in which individuals, influenced by others, engage in herding behavior, potentially distorting the decision-making process [Schöbel et al., 2016].

Another critical limitation of social information sharing, particularly evident in social networks, is the phenomenon of echo chambers [Cinelli et al., 2021]. In these echo chambers, individuals within like-minded groups limit their exposure to diverse perspectives, thereby fostering information bias. The echo chamber effect poses a challenge to the broad exploration of ideas and the inclusion of diverse viewpoints, potentially hindering the collective intelligence of the group.

1.5.2 Size and composition of the group

The concept of wisdom of crowds suggests that aggregating the judgments of many individuals can yield more accurate results than those of individual experts [de Condorcet, 1785; Galton, 1907; Surowiecki, 2005]. This phenomenon exploits the diversity of opinions and judgments within a group. Despite the possibility that individual estimates may deviate significantly from the true value, the aggregation of these diverse errors tends to cancel out, resulting in a more accurate collective judgment. Consequently, the larger the crowd, the greater the potential accuracy of the estimate. Moreover, empirical research on public goods games shows that group size can have a positive effect on cooperative behavior [Pereda et al., 2019].

However, the effectiveness of the wisdom of crowds depends on certain conditions, especially the independence or diversity of individual judgments [Page, 2008]. Experimental evidence supports this, highlighting that social influence can undermine the wisdom of crowds [Lorenz et al., 2011]. In cases where individuals observe each other’s beliefs, their estimates tend to converge, reducing diversity without a proportional increase in group accuracy. Additionally, when individuals are reluctant to change their beliefs beyond a certain threshold, this can lead to the formation of several distinct opinion clusters [Deffuant et al., 2000].

While the wisdom of crowds works well for simple situations, the complexity of real-world situations introduces nuances. Studies suggest that in these more complicated situations, smaller groups often outperform larger ones [Kao and Couzin, 2014]. This

underscores the importance of considering the nature and complexity of the task when determining the optimal group size for collective decision-making.

1.5.3 Network structure

The topology of social networks profoundly affects the transmission and assimilation of information within a group [Friedkin and Johnsen, 1999; Newman, 2003; Centola, 2022]. For example, studies have shown that groups organized as small-world networks exhibit enhanced signal-propagation speed and computational power due to a balance between local clustering and global connectivity [Watts and Strogatz, 1998; Barrat and Weigt, 2000]. Thus, understanding and intentionally manipulating network structures can be used to enhance collective decision-making [Horsevad et al., 2022].

In addition, Granovetter’s theory of the strength of weak ties [Granovetter, 1973] suggests that individuals who act as bridges between different clusters in a network facilitate the flow of information. These efficient networks thus facilitate information transfer and increase the collective problem-solving capacity of the entire group. However, Centola’s work on complex contagion [Centola, 2018] shows that counterintuitively, inefficient networks that facilitate knowledge transfer across groups can be more effective for complex problem-solving tasks.

1.5.4 The role of noise in group decision-making

Noise in the context of group decision-making can be conceptualized as extraneous and irrelevant information that introduces an element of unpredictability and randomness into the outcome [Tversky and Kahneman, 1974]. It materializes as a deviation from the relevant data and has a profound impact on the accuracy and coherence of group decisions.

It is important to distinguish between noise and bias. While bias involves systematic errors that consistently bias decisions in a particular direction, noise introduces a degree of variability and inconsistency that leads to erratic outcomes that can be difficult to predict and control [Kahneman et al., 2021].

To effectively mitigate the disruptive effects of noise, groups can adopt structured decision frameworks. For example, using redundancy in decision-making, where multiple people contribute to the decision or the same people evaluate the decision multiple times, can help identify and counteract the effects of noise. Another increasingly common approach is to use algorithms to facilitate decision-making. By systematically processing information and reaching consistent decisions, algorithms have the potential to reduce noise.

However, it is important to recognize that algorithmic noise reduction is not without its challenges. While algorithms can reduce noise, they can also introduce bias and discrimination into the decision-making process. This is particularly difficult to identify and correct because the data used to train algorithms may have inherent biases, and correlations within the data can perpetuate and exacerbate existing biases [O’Neil, 2017].

Achieving the optimal noise balance to maximize collective performance is a nuanced task. While aiming for zero noise may seem practical, it may not be desirable. Indeed,

1.6. Outline of the Thesis

a low level of noise can contribute to a diversity of perspectives and prevent groupthink, fostering a more robust decision-making environment. Therefore, the focus should be on managing and calibrating the level of noise to improve decision outcomes without compromising the integrity of the process.

1.5.5 Social nudges and bots

Nudges, which are subtle choice designs that guide individuals toward certain choices without prohibiting options or restricting individuals, are becoming more common in our society [Thaler and Sunstein, 2008]. Nudging has been used as a tool with more or less success [Kosters and Van der Heijden, 2015] by governments and companies [Moseley and Stoker, 2013]. In addition, online recommendation systems use digital nudging to influence user behavior through personalized suggestions, influencing choices in areas such as e-commerce and social media [Jesse and Jannach, 2021].

Recognizing the impact of nudges on individual decision-making, there is potential to harness these principles to enhance collective intelligence. By strategically implementing nudging strategies in collaborative settings, one may optimize group dynamics and foster more informed and effective collective decision-making processes [Mele et al., 2021].

One of the possible methods to nudge individuals is to use social bots, which are automated agents on social platforms. These tools have the potential to enhance collective intelligence [Jayles et al., 2020a; Ali Mehenni et al., 2021]. However, bots on social platforms are often used maliciously to spread misinformation [Ferrara et al., 2016; Broniatowski et al., 2018]. Therefore, ethical considerations regarding transparency and accountability need to be carefully addressed.

1.6 Outline of the Thesis

In this context, it is clear that digital traces are playing an increasingly important role in shaping our daily lives and influencing human decision-making [Chavalarias, 2022]. Given their susceptibility to manipulation, it becomes imperative to gain a deeper understanding of their societal impact. Therefore, the main objective of this Ph.D. thesis is to examine how humans use and interact with digital traces, to improve our understanding of their effects on human decision-making and society as a whole. The research will specifically analyze the conditions under which cooperation can arise through the use of digital traces, and examine how competition influences human utilization and interaction with these traces.

To address these questions, the thesis begins with three chapters devoted to materials and methods. In Chapter 2, I present the STIGMER game, a game developed before the start of the thesis and refined during the thesis. This game, designed as an information search task inspired by recommendation systems, serves as the basis for all analyses in the thesis. In Chapter 3, I present the main observables essential to our analysis, justifying their selection and explaining their usefulness. Chapter 4 outlines the principles of the agent-based model used to model the behavior of a player of the STIGMER game. This versatile model allows us not only to create agents that mimic human behavior for a

deeper understanding of their behavior but also to identify optimal strategies in different situations. In addition, agents following this model can play alongside humans and actively interact with them.

The thesis then proceeds through three result-oriented chapters. Chapter 5, elucidates the ability of humans to use stigmergic interactions to promote cooperation within a group. It also shows the effect of competition on the behavior of participants and its impact on cooperation and deception. Chapter 6 examines the effect of group composition on individual and collective performance, using a setup where a human interacts with four bots governed by the model described in Chapter 4. Chapter 7 examines the impact of different levels and types of competition on individual and collective behavior.

Finally, Chapter 8, serves as the thesis conclusion. This is where the synthesis of the findings takes place, with a discussion that reinforces the implications of the results in the broader context of the research landscape.

Experimental Design

Chapter Summary

The STIGMER game is a multiplayer game in which participants engage in a collective search for information in a digital environment. Originally designed by Guy Theraulaz and Clément Sire, several extensions of the STIGMER game have been developed and studied experimentally and theoretically during my Ph.D. thesis. The STIGMER game reproduces certain aspects of rating systems used on many platforms on the Internet. The primary objective for the players is the identification of the best options among available ones, with the help of only indirect interactions in the form of collective ratings.

In the game, participants are presented with a grid of cells (the available options) each containing a hidden value. In successive rounds, each player sequentially discovers and rates N_{turns} different cells. After each round, the colors of the cells change according to the ratings given by the players since the beginning of the game. The colors of the cells, based on these ratings, are the only means of communication available to the participants.

The game is packaged in a versatile web application that can be accessed from any location and device. This online platform provides the flexibility to run multiple game variations, each with subtle changes in the gameplay such as number of players, round, score calculation, etc. It provides participants with a user-friendly interface for an enjoyable gaming experience and gives us convenient access to configure and track the progress of experiments.

During my Ph.D. studies, we conducted several series of experiments, which are summarized at the end of the present chapter and addressed in detail in Chapters 5 to 7. Each session lasted one hour, during which the participants played a dozen games. At the end of the session, participants received a monetary reward that could be correlated to their performance in finding the best options in the game. Changing this reward changes the incentives of the players and allows to tune the intragroup and intergroup competition. Therefore, this experimental setup is an excellent tool for studying cooperation and competition processes within and between groups.

The anonymized data collected from all of these experiments serve as the basis for the analysis and modeling in Chapters 5 to 7.

Contents

2.1	The Stigmer game	18
2.1.1	Grid and cell values	19
2.1.2	Round	19
2.1.3	Colors	20
2.1.4	Score	23
2.2	Implementation of the Game	24
2.2.1	Web application	24
2.2.2	Visualization application	27
2.3	Experimental Setup	27
2.3.1	Experimental session	27
2.3.2	Payment	30
2.3.3	Data	30
2.3.4	Participants	31
2.3.5	Ethics statement	31
2.4	Summary of the Conducted Experiments	32

The research presented in this thesis originated from a collaboration between the *Laboratoire de Physique Théorique* (LPT), the *Centre de Recherches sur la Cognition Animale* (CRCA), and the *Toulouse School of Economics* (TSE). The main objective was to investigate stigmergic interactions within human groups. To achieve this, a novel game was developed, inspired by the rating systems used in many online commercial websites, in which participants engage in a collective information search task through a digital interface. In this game, players must find the best available options in their environment by solely communicating via indirect interactions taking the form of digital traces.

This chapter begins with an introduction to the theoretical basis of the game, followed by a description of its implementation. Then, I present the experimental procedure, detail the data collection, and briefly describe several variations of the game (and their main motivations) used in the actual experiments, which are analyzed in Chapters 5 to 7.

2.1 The Stigmer game

The idea behind the game, which we named STIGMER, is to reproduce processes observed in a star rating system [Lü et al., 2012]. In this game, N_{players} players are instructed to identify the cells with high values within a grid containing hidden numbers. The only communication form allowed among them is the ratings they assign to the cells they explore.

This game draws parallels to popular online marketplaces and platforms like Amazon, TripAdvisor, and eBay, where users rate products, services, or sellers, aiming at finding the best possible choice. While we acknowledge that the game does not capture all aspects

2.1. The Stigmer game

44	99	0	28	3	0	4	2	1	9	1	0	8	11	1
2	4	44	2	1	1	3	4	2	7	4	2	3	44	3
8	2	3	4	2	2	1	1	2	0	1	3	2	45	0
3	6	2	1	6	2	0	71	3	8	2	3	4	20	20
0	1	46	85	0	2	43	3	1	1	1	1	6	5	2
2	2	13	1	0	0	24	86	0	3	2	2	1	3	20
1	1	20	4	1	2	13	21	22	0	2	2	0	1	6
3	4	13	1	3	2	1	1	24	6	53	3	3	2	19
5	3	84	3	2	0	0	3	2	27	3	1	3	14	2
86	3	2	6	1	45	20	7	24	2	4	27	3	2	3
1	1	7	0	43	7	1	0	4	2	1	1	53	3	0
72	2	2	2	6	8	19	3	1	1	72	1	1	1	14
3	2	4	2	7	0	0	12	0	71	1	3	53	1	9
2	2	1	8	2	1	0	12	1	2	3	2	1	1	1
3	1	2	27	51	2	2	3	2	0	2	0	2	3	6

Figure 2.1: **Example of a grid with the corresponding cell values.** Random shuffling of the 225 values of the 15×15 grid used in the game.

of real-life situations, it does investigate analogous processes. These processes include: (i) the exploration of available options (cells in our game; products in an online store), which are strongly influenced by their current ratings, (ii) the participants' evaluation of these options, which is primarily influenced by the quality of the chosen option, and (iii) the dynamical evolution of the ratings.

2.1.1 Grid and cell values

Players face a grid consisting of a total of N cells, arranged in a grid. Each cell contains a hidden value. In our experiments, we used a 15×15 grid, so that $N = 225$. An example of such a grid is shown in Figure 2.1.

The values within the cells are integers ranging from 0 to 99. The distribution of these values exhibits a predominance of small values, with only a few larger ones. The mean value across all cells is 9.7, while the median value is only 2. The exact distribution of these values is shown in Figure 2.2. Due to the scarcity of high-value cells, this distribution should encourage cooperation if agents want to uncover these high-value cells.

It is important to note that although the values remain the same throughout each game, they are initially shuffled, resulting in a unique arrangement for every game played.

2.1.2 Round

A game is divided into N_{rounds} successive rounds. In each round, each player must visit and rate N_{turns} distinct cells, independently of the others. In the actual experiments, we use $N_{\text{rounds}} = 20$ and $N_{\text{turns}} = 3$, which offers a compromise between a large enough number of visited cells ($3 \times 20 = 60$ per player for a total of $300 > N = 225$ cells for




Figure 2.2: **Distribution of the cell values.** Distribution of the 225 cell values used in the game. The mean is 9.7 and the median 2.

five-players games) and a short enough game duration allowing to conduct a dozen games per one-hour session.

When a cell is visited, its hidden value is revealed to the player (see Figure 2.3). The player is then required to rate the cell using a scale of 0 to 5 stars. After rating the cell, the value disappears, and a cross indicates that the cell has been visited. The player can then proceed to visit and rate another cell. If the player has already visited and rated N_{turns} cells, they must wait for the other players to do so. It is important to note that players are not permitted to revisit the same cell more than once within a single round. However, they may revisit cells in subsequent rounds.

After each player has visited and rated their N_{turns} cells, the round ends. The ratings from all players are then combined, and the colors of the cells are updated to reflect the collective ratings. The players then proceed to the next round and repeat the same process.

2.1.3 Colors

Players do not have direct access to each other’s ratings. Instead, they perceive the ratings through the colors of the cells, which change over time. The colors are related to the fraction of stars each cell has received since the start of the game, which is calculated as the sum of all ratings for that particular cell divided by the sum of all ratings for all cells.

The colors range from white to black through a gradient of red (see Figure 2.4). In this gradient, white represents the absence of stars, indicating that a cell has either *never been visited* or *has been visited but always rated 0*. Conversely, black indicates that the cell contains all the stars. The intensity of the red hue becomes darker as the fraction of stars within the cell increases. See Figure 2.5 for an example of a game where the evolution of the colors of the cells can be observed during the game.

The exact RGB color value for a cell, $c(P)$, corresponding to the fraction of stars it

2.1. The Stigmer game




Figure 2.3: **Illustration of the cell selection and rating.** The image displays the grid of hidden values, with two black crosses indicating the cells that have already been visited during the round. Additionally, it shows the value of the currently selected cell and the rating selection interface.

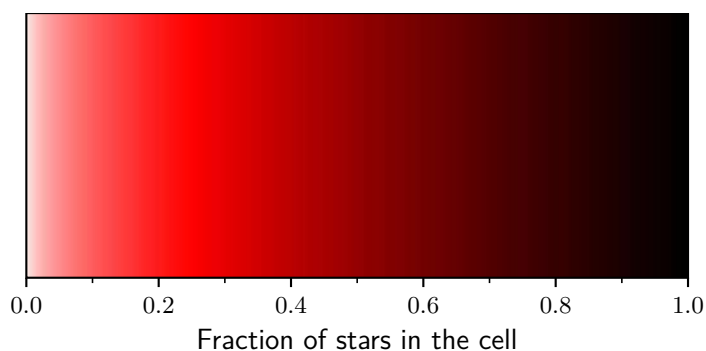


Figure 2.4: **Color scale.** Color scale of the visited cells as a function of the fraction of stars used to rate the cells since the beginning of a game.

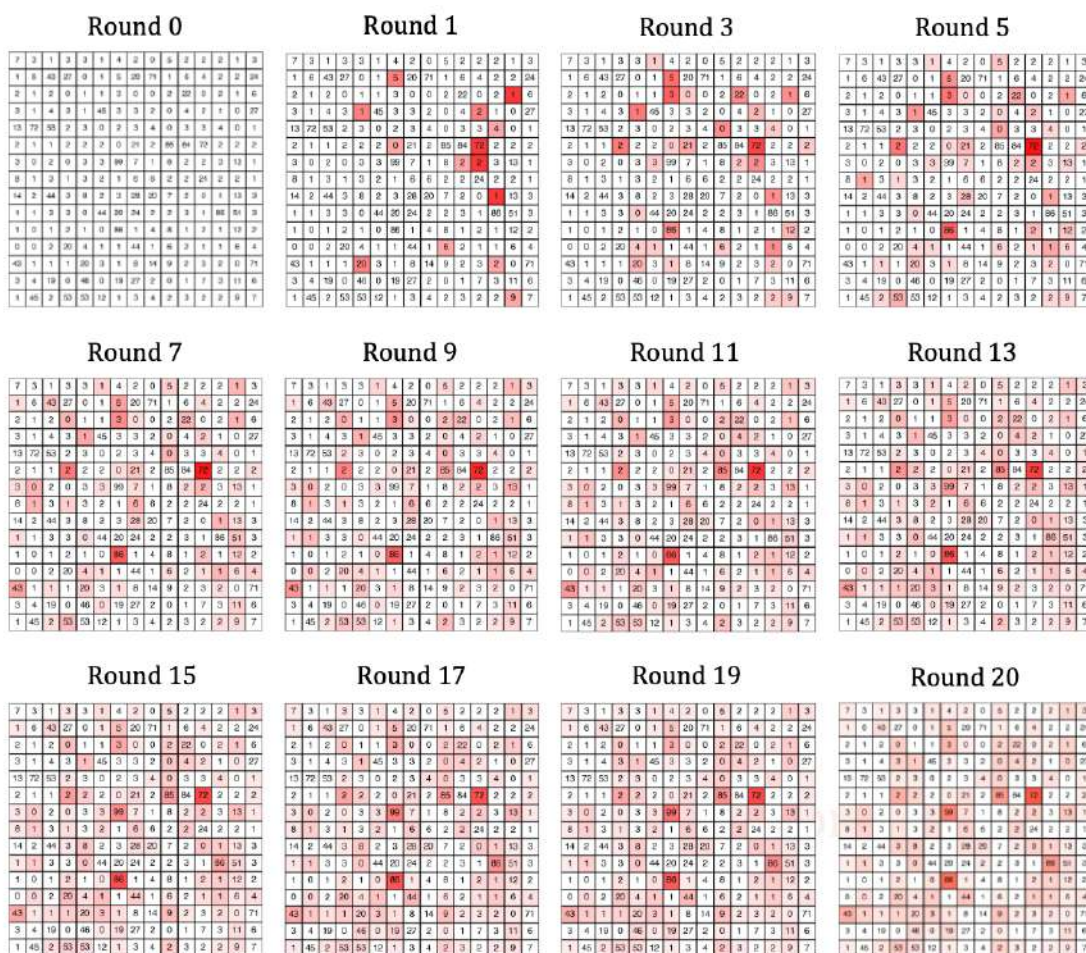


Figure 2.5: Evolution of the colors of the cells during the game. At the beginning of the game (round 0), all cells are white. As the round number increases, more cells are colored, with some cells being darker than others. All cell values are displayed in this figure, but they were hidden from the players during the game: players could only see the value of the cells that they were currently visiting.

2.1. The Stigmer game

contains, P , is determined using the following formula:

$$c(P) = \begin{cases} \text{RGB}(1, 1 - 2x(P), 1 - 2x(P)) & \text{if } 0 \leq x(P) \leq 1/2 \\ \text{RGB}(2(1 - x(P)), 0, 0) & \text{if } 1/2 < x(P) \leq 1 \end{cases} \quad (2.1)$$

where $x(P) = \sqrt{P}$. Note that the color does not scale linearly with the fraction of stars, but rather with the square root of the fraction. This decision was made because of the typical small magnitude of the fractions P involved. A linear scale $x(P) = P$ would not provide enough contrast, resulting in most cells appearing almost white. The square root, on the other hand, provides better visibility, ensuring that cells are clearly distinguishable.

At the start of the game, the fraction of stars within each cell being undefined, a convention has been adopted where every cell is white. This convention ensures that the white color uniformly represents the absence of stars in each cell. With this convention, the game maintains a coherent visual representation, allowing players to easily identify and differentiate cells that have yet to receive any stars.

2.1.4 Score

The game offers the option to play with or without a scoring system. In the former scenario, the player's score for each round is determined by their actions during that round. This includes the value of the cells visited by the subject and, if applicable, their corresponding rating.

In all experiments presented in this work (see a summary in Section 2.4 of this chapter and a more detailed description in Chapters 5 to 7) where the game included a score, the score of a player increases by the value of each visited cell. Hence, the score is *independent of the rating* given by the player to the visited cells. The final score of a player is calculated by summing the values of the cells visited during the twenty rounds of the game. One can then even define the score of a group of players playing on the same grid as the sum of the final scores of the group members.

Note that other definitions of the score, not addressed in this Ph.D. thesis, can be investigated, illustrating the richness of the STIGMER game. For instance, a player's score may increase by the value of the cells visited multiplied by their rating. To prevent having only five-star ratings, there needs to be a maximum number of stars (eight, for example) that can be used in a given round. This score definition encourages players to carefully choose which cells to place the eight stars and allocate them to the cells with the highest values. It is a form of nudging that encourages players to mark high-value cells, which appear darker on the grid and disclose them to other players. The score definition fosters collaboration implicitly.

When the game includes a scoring mechanism, participants receive real-time updates on their cumulative scores throughout the game. This feature allows them to actively track their progress, giving them valuable insights into their performance. Furthermore, upon reaching the end of the game, players are presented with a comprehensive table that displays the final scores and corresponding rankings of all players in the game. This table serves as a comparative reference, allowing each player to evaluate their performance

within the game relative to that of the others. By presenting both the final score and ranking, players can gain a comprehensive understanding of their achievements and overall standing within the game. This comes important when considering repetitions of this game.

2.2 Implementation of the Game

2.2.1 Web application

The experiments were conducted using a web application developed specifically for the STIGMER game.

A first version of the application was developed by the company ANDIL (<https://www.andil.fr>) for the preliminary experiments conducted in 2017. This version ran on an admin computer, with player computers connected to the admin computer through a Local Area Network (LAN). The first version solely included a user interface for the players, while the admin had to manually modify the code of the application to adjust game parameters.


During my Ph.D., a second version of the application was developed. This version was primarily built by Maxime Delacroix, a full-stack developer intern at the Laboratoire de Physique Théorique, with me overseeing and assisting him in the scientific aspects. This second version introduced numerous improvements compared to the initial one. This second version has been developed in JavaScript using Node.js and MongoDB for database integration.

While the player user interface (see Figure 2.6) underwent only minor changes, such as the option to switch between French and English, the notable addition was the introduction of an admin user interface (see Figure 2.7). This new interface allowed the creation of new games with new rules directly from the web application (see Figure 2.7A), which greatly simplifies the game creation process. Additionally, it incorporated tools for real-time visualization of game progress and scores (see Figure 2.7B). Another notable enhancement of this second version of the web application was the implementation of bots, allowing games to be played with a combination of real players and computer-controlled bots.

The updated version of the application has been deployed on the server of the Laboratoire de Physique Théorique and made available online at the following URL <https://www.lpt.ups-tlse.fr/Stigmer/local>. This online availability has two advantages. First, it greatly facilitates in-person experiments, since participants only need an Internet-connected device to participate, eliminating the need for a local network connection. This streamlined accessibility simplifies the logistics of conducting experiments in physical settings. Second, online availability opens up the possibility of conducting experiments autonomously, without the need for direct human supervision. Although our current experiments are exclusively in-person, we anticipate experimenting online in the future. This could involve sharing the game link on social networks or using crowdsourcing platforms such as Amazon Mechanical Turk to extend our research reach and engage participants in an online context.

2.2. Implementation of the Game

A



B



Figure 2.6: **Player interface of the web application.** (A) Player interface during the game. The remaining time, the current score, and the current round are displayed at the top. In the center of the interface is the grid with the colored cells, the black crosses marking the cells that have already been visited during the round, and the value of the currently selected cell along with the rating selection interface. At the bottom is a list of all the players in the game. (B) Player interface at the end of the game, displaying the final scores of the players in the group.



Figure 2.7: **Admin interface of the web application.** (A) Admin interface to create a new type of game. (B) Admin interface to visualize the live progression of the games.

2.3. Experimental Setup

For a deeper understanding of the game, an online version is available at the following URL <https://www.lpt.ups-tlse.fr/Stigmer/web>. This version allows anyone to test their skills against four computer-controlled bots.

The application is open source, and the most recent version can be accessed on a GitHub repository at the following link: <https://github.com/Thomas-bssnt/Stigmer-web-app>.

2.2.2 Visualization application

I also developed another application to visualize the games. This tool allows a detailed round-by-round visualization of the stars and visits distributions. The application was implemented in Python, utilizing the Tkinter library for the graphical user interface. An example of visualization is shown in Figure 2.8.

This visualization tool is open source and available in a dedicated GitHub repository at the following URL <https://github.com/Thomas-bssnt/Stigmer-visualization-app>.

2.3 Experimental Setup

A first preliminary/exploratory series of experiments took place in 2017, before the start of my Ph.D., followed by four additional series between December 2021 and January 2023, during my Ph.D.

Each series of experiments lasted five or six days. Within each experimental series, several sessions (five or six per day), were organized in which participants came and played typically a dozen games for a total of one hour. Within a given experimental series, the games played were not necessarily the same from one session to the next, and sometimes differed in their incentives and payment structure.

The different experiments conducted during my Ph.D. thesis are summarized in Section 2.4 of this chapter and analyzed in Chapters 5 to 7.

Three of these four series were conducted at the Toulouse School of Economics (TSE), while the experiments with bots were conducted at the Laboratoire de Physique Théorique (LPT). Indeed, the participants were playing against bots without their knowledge but were informed that they were playing against each other. Deceptive experimental designs are generally prohibited by institutions such as TSE and economics journals, but not in social psychology. Due to this restriction, our collaboration organized these experiments with bots at the LPT, under the supervision of the ethical committee of the Université de Toulouse.

2.3.1 Experimental session

The participants were requested to arrive ten minutes before the beginning of the experimental session. During this period, they waited in a designated waiting room (see Figure 2.9A), where they read the information note, and filled and signed the informed consent form (see Appendix A.1). They also randomly selected a player ID, which deter-

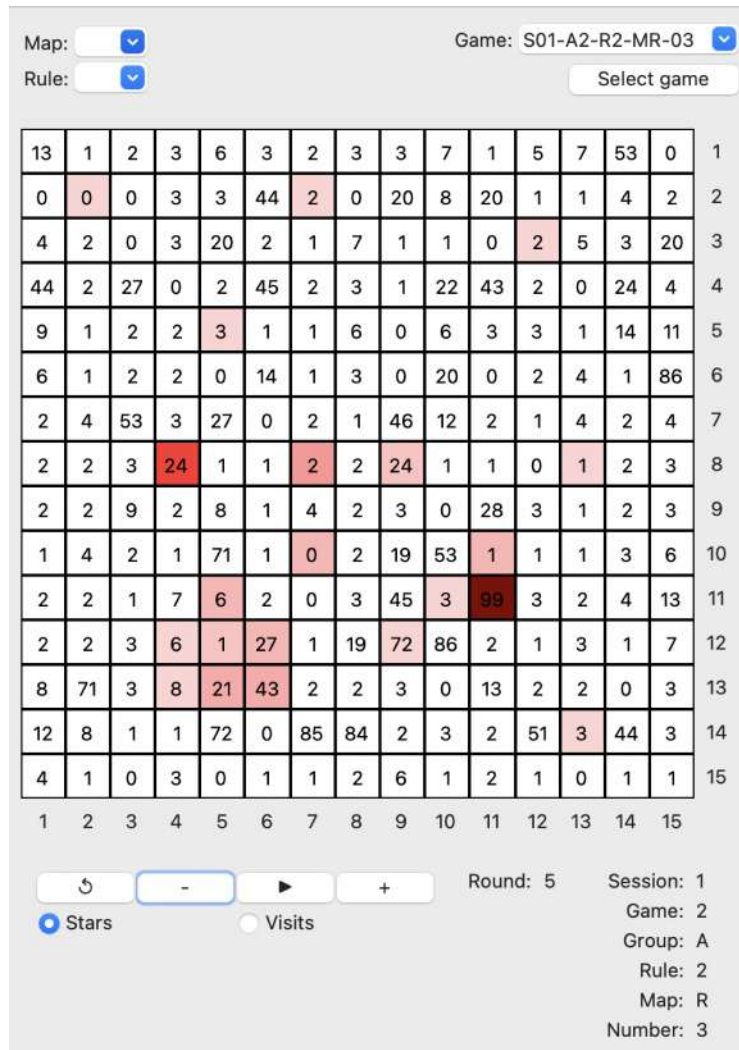


Figure 2.8: **Interface of the visualization application.** Example of game visualization in the visualization application, with the game selection area at the top, the control and information area at the bottom, and the visualization in the middle.

2.3. Experimental Setup



Figure 2.9: **Photographs of the experiments.** Photographs of the experiments conducted at the Behavioral Economics Laboratory of the Toulouse School of Economics. (A) Waiting room. (B) Room in which the rules were explained to the participants. (C) Experimental room. (D and E) Participants playing the STIGMER game. (F) Payment of a participant.

mined their location in the experimental room and their group if more than one group was playing at the same time.

Once all participants were present, they were led into another room where the rules of the game, payment procedures, and guarantees of anonymity were explained (see Figure 2.9B). While participants were encouraged to ask questions about the rules, questions about strategy were not answered.

Following the briefing, participants proceeded to the experimental room, each occupying a cubicle corresponding to their randomly selected player ID (see Figures 2.9C and 2.9D). This arrangement ensured no interactions among participants. To guarantee a comprehensive understanding of the game, a recap of the basic rules was reiterated before starting the first game.

Then the experimental phase began (see Figures 2.9D and 2.9E). From my computer, I assigned participants to games and started them. The duration of each game was typically three to four minutes. As soon as one game was finished, another was started until the desired number of games (about twelve per session) was reached. To respect the desired number of games, a soft time limit of twenty seconds was set for each round. Beyond that, a message was displayed at the top of the player's screen encouraging late players to play faster. This limit was largely sufficient for the vast majority of participants, but in each session, a few were always slow. If no acceleration was observed after a few games, we went to them to understand what was causing the delay and gently asked them to speed up. On the rare occasions when the web application crashed, or when someone was slow

and did not accelerate, participants were stopped before reaching the required number of games to avoid delays in subsequent experimental sessions.

Notably, approximately 20% of the registered participants either did not come or arrived considerably late. Since a specific number of players was required to start an experimental session, substitute participants had to be found. These substitutes were recruited from the building, directly from the street, or occasionally from previous experimental sessions. During the last experimental series, we booked one or two additional participants for each session to ensure an adequate number of participants. If more participants than necessary arrived on time, we provided monetary compensation to those who were present but did not participate in the experiments.

At the end of the experimental session, participants were paid and signed a receipt (see Figure 2.9F). For the last two experimental series, the participants also filled out a questionnaire (see Appendix A.2) on their understanding of the game and their strategies.

2.3.2 Payment

At the end of each session, all participants were paid, with an amount that may vary for each individual. Assuming that participants aim to maximize their rewards, modifying the payment scheme is expected to influence participants' behaviors.

In most cases, the payment was based on the cumulative score of the participants throughout all the games played within the session. However, we also explored variations in the payment scheme where participants were paid according to the score of the group, and others where every participant was paid the same.

For instance, in Chapter 5, we will explore the consequences of two payment structures: one in which all participants receive equal payment, regardless of their performance, and another in which participants are rewarded based on their ranking within their group.

In the experiments conducted at the Toulouse School of Economics, participants received cash payments (see Figure 2.10A), while in those performed at the Laboratoire de Physique Théorique, participants were compensated with Amazon gift cards of various amounts (see Figure 2.10B).

2.3.3 Data

For each game, the web application systematically archived a variety of data in a server-side database. This included information such as details about the session itself, the characteristics of the game (including the number of rounds, players, and rule parameters), and the demographics of the participants, including age and gender. Crucially, it also recorded the actions performed by the players during the game, including their visits and ratings.

At the end of each session, this data could be conveniently exported into a ZIP file containing two different files for each game. The first file, in JSON format (see Appendix A.3 for an example), contains information about the session, the players involved, and the exact configuration of the game. This JSON file allowed both identification and reproduction of the game. In addition, a second file in CSV format (see Appendix A.4 for an

2.3. Experimental Setup




Figure 2.10: **Payment given to the participants.** (A) Cash. (B) Amazon gift cards.

example) was created to specifically contain detailed records of player actions during the game.

Complementing this data was the questionnaire (see Appendix A.2). This questionnaire, which was distributed at the end of each experimental session, can be used as a complementary tool to gain insight into the participants' experiences and strategic choices during the game.

2.3.4 Participants

Participants were recruited through email advertisements distributed to various research labs and academic institutions in Toulouse, including the Toulouse School of Economics (TSE), TBS Education, Université Toulouse Capitole, and Université Toulouse III – Paul Sabatier. In addition, recruitment efforts were supported by leveraging the TSE participant pool, which includes both students and other individuals. As a result, the majority of participants in our experiments were students.

A total of 1005 participants took part in the experiments. The age and gender distribution is presented in Figure 2.11. The data indicates a slightly higher number of female participants than male participants. The majority of participants were between the ages of 18 (the minimum age) and 26, with fewer individuals older than this range.

2.3.5 Ethics statement

The aims and procedures of each experiment received ethical approval from one of two ethics committees (see Appendix A.5). The three series of experiments conducted at the Toulouse School of Economics were approved by the Toulouse School of Economics Research Ethics Committee for Experimental Research. Similarly, the experiments involving bots conducted at the Laboratoire de Physique Théorique were approved by the Comité d'Éthique de la Recherche of Université de Toulouse.




Figure 2.11: **Participant demographics.** Age and gender distribution of the 1005 participants in our experiments. (A) Age distribution on a logarithmic scale, where the mean age is 22.22 and the median age is 21. (B) Gender distribution, consisting of 46 % male participants and 54 % female participants.

2.4 Summary of the Conducted Experiments

This section provides a brief overview of the experiments and their motivations that will be analyzed in the subsequent chapters. A more in-depth description will be presented at the beginning of each chapter dedicated to the analysis of these experiments.

In Chapter 5, we study the impact of competition within a group on cooperation in two series of experiments. In the first experiment there is no notion of score and the five participants in a group are paid 10€ each. Thus, there is no competition among the participants and their behavior does not influence their reward. In the second experiment, the five participants in a group are ranked based on their scores and paid accordingly. The participant ranked first is paid 20€, the second is paid 15€, and the three remaining participants are paid 10€ each. In this experiment, the behavior of the participants in the game directly impacts their reward, and there is a competition among the group members. The methodology and analysis of this Chapter 5 follow the lines of [Bassanetti et al. \[2023\]](#).

In Chapter 6, we study the impact of the group composition on cooperation. We performed experiments with five participants per session, but instead of playing together in the same group, each participant actually played against four bots commanded by our model, unbeknownst to them. At the end of the experimental session, the five human participants are ranked together based on their score and paid accordingly. The first is paid 20€, the second 15€, and the three remaining participants are paid 10€ each. By controlling the behavior of the bots, the setup provided a control of the cooperation level within the group. We performed nine different variations of this experiment in which the bots adopted different behaviors (collaborative, deceptive, or neutral). In Chapter 6, we also compare the results of these experiments to the competitive experiment analyzed in Chapter 5, in which five humans played together. This work is the subject of an article about to be submitted [[Bassanetti et al., 2024a](#)].

In Chapter 7, we study the impact of intragroup and intergroup competition on individuals' behaviors and cooperation. This work is the subject of an article in preparation

2.4. Summary of the Conducted Experiments

[Bassanetti et al., 2024b]. We compare seven experiments where the participants play in groups of five, which can be classified into four categories depending on the type of competition: no competition (2 experiments), intragroup competition only (1 experiment), intergroup competition only (1 experiment), and a mix of intragroup and intergroup competition (3 experiments):

- There are two experiments with no competition directly induced by the payment structure. The first experiment does not have any notion of score and is analyzed in detail in Chapter 5 and in Bassanetti et al. [2023]. In the second experiment, participants are paid proportionally to their score, so that their final rank in their group is irrelevant.
- The experiment with pure intragroup competition is studied in detail in Chapter 5. The participant ranked first is paid 20€, the second is paid 15€, and the three last participants are paid 10€ each.
- In the experiment with intergroup competition only, two groups of five participants are directly competing. The five members of the group with the highest group score are paid 20€ each, and the members of the other group are paid 10€ each.
- There are three experiments with both intragroup and intergroup competition involving ten players divided into two groups of five. At the end of the experimental session, the ten players are ranked together based on their score and paid accordingly. These three experiments differ by their degree of intragroup competition compared to the intergroup competition, which is tuned by the payment structure. In a first experiment, the participants ranked first through fifth are paid 15€ each, and the other ones ranked sixth through tenth received 8€ each. In a second experiment, the participant ranked first is paid 15€, the second 14€, the third 13€, and so on, down to the last participant, who is paid 6€. Finally, in a third experiment, the participant ranked first is paid 20€, the second and third 15€ each, and the seven last participants are paid 10€ each (this experiment is analyzed in detail in Bassanetti et al. [2023]).

Main Observables

Chapter Summary

This chapter serves as an essential foundation by introducing key notations and observables that are used in the subsequent chapters.

First, it presents the observables that are used to quantify collective behavior, including the normalized group score for group comparisons, visit, and rating performance to assess players' exposure to high-value cells, the inverse participation ratio, which quantifies the number of cells visited and rated by the players during the game, and the fidelity, which captures the correlation between grid values and players' visits and ratings.

The chapter then describes the observables used to measure individual visit behaviors, including the normalized player score for comparison between players, the best values of cells opened by a player, the probability of revisiting these best cells, and the probability of finding high-value cells, providing insight into cooperative tendencies.

We then describe the observables used to quantify individual rating behaviors, primarily the mean rating given to a cell, which is the main tool used to classify individual behaviors, and the fractions of ratings, which provide further information about rating behaviors.

Finally, the chapter introduces the bootstrap method, which is the tool used to calculate error bars in the following chapters.

Contents

3.1	Notations	36
3.2	Quantifying the Collective Behavior	37
3.2.1	Normalized group score	37
3.2.2	Fractions of visits and fractions of stars	38
3.2.3	Visit and rating performances	39
3.2.4	Inverse participation ratio	40
3.2.5	Fidelity	41
3.3	Quantifying Individual Visit Behaviors	42
3.3.1	Normalized individual score	42
3.3.2	Values of visited cells	43
3.3.3	Probability of revisiting cells from the previous round	44
3.3.4	Probability of finding cells with the highest values	44
3.4	Quantifying Individual Rating Behaviors	45
3.4.1	Mean rating	45
3.4.2	Fitting of the mean rating	46
3.4.3	Fractions of ratings	47
3.5	Computation of the Error Bars: the Bootstrap Method	47

This section serves as an introduction to the key observables that will be used to quantify both the collective and individual behaviors examined in each chapter of this thesis. We provide definitions and explanations of these observables, establishing their significance and relevance to our study.

3.1 Notations

Let us first introduce the notations that will be used throughout this article.

Vectors are finite collections of components. They are represented in this thesis by bold letters. For example, \mathbf{x} denotes a vector. The individual components within these vectors are indicated by subscripts, such that x_i represents the i -th component of \mathbf{x} . The scalar product of two vectors \mathbf{x} and \mathbf{y} , denoted $\langle \mathbf{x} | \mathbf{y} \rangle$, is defined as $\langle \mathbf{x} | \mathbf{y} \rangle = \sum_i x_i y_i$. Similarly, the Euclidean norm of a vector \mathbf{x} , denoted $\|\mathbf{x}\|_2$, is defined as $\|\mathbf{x}\|_2 = \sqrt{\sum_i x_i^2}$.

Continuing with some letter notations used consistently for ease of understanding:

- t represents the round of a game, it is the equivalent of time in our experiment. It is an integer between 0 and N_{rounds} .
- V represents the value of a cell. It is an integer between 0 and $V_{\text{max}_1} = 99$.
- c represents the index of a cell. The two-dimensional grid of 15×15 cells is flattened to a one-dimensional array of 225 cells. Thus, c is an integer between 1 and $N = 255$.

3.2. Quantifying the Collective Behavior

- s represents a rating, where “s” stands for “stars”. It is an integer between 0 and 5.
- $\langle O \rangle$ represents the mean of an observable O (defined in the next sections) averaged over all games played for a given experimental condition (incentive rule, payment structure, participants playing against bots...). Indeed, for each experimental condition considered in this thesis, we have conducted typically 100-200 games (see details when we address a specific condition), and the mean of an observable is understood as the mean over all such games. For individual observables, $\langle O \rangle$ is the average of O over all played games, and over all players of these games.

The vector containing the hidden values of the grid will be denoted by \mathbf{V} , where V_c corresponds to the value of the cell with index c . The highest, second-highest, and third-highest values within the grid will be respectively referred to as V_{\max_1} , V_{\max_2} , and V_{\max_3} .

Note that the letters “P” and “Q” are used for many observables, so it is important not to mix them up. For example, $p_c(t)$, $\mathbf{p}(t)$, $p(t)$, $P_c(t)$, $\mathbf{P}(t)$ and $P(t)$ are all different quantities

3.2 Quantifying the Collective Behavior

Let us now delve into the observables used to quantify the collective behavior of a group. These observables are defined based on the actions of all participants, without distinguishing between them.

3.2.1 Normalized group score

Section 2.1.4 introduced the concept of score in a game. Note that games can be played with or without scoring, and in cases where scoring is used, there may be different types of scoring. To allow for cross-game comparisons, we introduce a uniform scoring metric in our analysis, which is applied consistently to all games, regardless of their specific scoring method.

This newly proposed final score is defined as the cumulative sum of the values discovered by the players throughout the game. For a more meaningful and easier interpretation, this score is then normalized by the maximum possible score, which depends on the grid values and the number of game rounds.

Definition 3.1 — Normalized group score

The normalized score of a group, denoted \hat{S} , is defined as the sum of the values visited by every player of a group during the game divided by the maximum possible score of the group. This maximum score corresponds to the score the group would achieve if every player had selected the three best cells in each round, resulting in a total value of $(V_{\max_1} + V_{\max_2} + V_{\max_3}) \times N_{\text{rounds}} \times N_{\text{players}}$, where V_{\max_1} , V_{\max_2} , and V_{\max_3} are the three highest value cells, and N_{rounds} and N_{players} are the number of rounds and the number of players, respectively.

The group score is normalized to a range between 0 and 1. A score of 1 corresponds to the best theoretical result, while a score of 0 corresponds to the worst possible result, indicating

that players only opened cells with a value of 0. The normalized group score serves as a metric for evaluating the collective performance of the entire group and facilitates comparisons between different groups.

An important quantity derived from the normalized group score is the mean normalized group score, that is, the score averaged over the games played in a given experimental condition.

Definition 3.2 — Mean normalized group score

The mean normalized group score, denoted as $\langle \hat{S} \rangle$, is defined as the arithmetic average of the normalized group scores over all games for a given experimental condition.

This mean normalized score is extremely useful for comparing experiments and quickly assessing the impact of rule changes or incentive adjustments on performance.

3.2.2 Fractions of visits and fractions of stars

In order to define the rest of the observables that describe the collective behaviors, we introduce the following four fundamental vectors.

First, we introduce the vectors characterizing the instantaneous and cumulative fractions of visits at round t .

Definition 3.3 — Fractions of visits

For each round t , we define the following vectors:

- The **instantaneous fractions of visits vector**, denoted as $\mathbf{q}(t)$, is the vector whose component $q_c(t)$ of index c represents the fraction of visits by all players to the cell c , at round t .
- The **cumulative fractions of visits vector**, denoted as $\mathbf{Q}(t)$, is the vector whose component $Q_c(t)$ of index c is the fraction of cumulative visits by all players to the cell c , from the first round up to and including round t .

Likewise, we introduce similar quantities for ratings instead of visits.

Definition 3.4 — Fractions of stars

For each round t , we define the following vectors:

- The **instantaneous fractions of stars vector**, denoted as $\mathbf{p}(t)$, is the vector whose component $p_c(t)$ of index c is the fraction of stars deposited by all players in the cell c , at round t .
- The **cumulative fractions of stars vector**, denoted as $\mathbf{P}(t)$, is the vector whose component $P_c(t)$ of index c is the fraction of stars deposited by all players in the cell c , from the first round up to and including round t .

Note that these two vectors are undefined if no stars have been put at round t for $\mathbf{p}(t)$ or up to round t for $\mathbf{P}(t)$.

These four quantities can be viewed as probability vectors, which means that their

3.2. Quantifying the Collective Behavior

components lie in the range $[0, 1]$ and that the sum of their components is equal to 1.

While these vectors provide valuable insights, the challenge lies in their high dimensionality ($N = 225$ components in our case), which requires new observables that can capture the complexity in a more comprehensive way.

3.2.3 Visit and rating performances

In the experiments, participants were explicitly asked to find cells with high values, so it is interesting to quantify how well they achieve this goal. To this end, we introduce four new quantities that quantify their performance in finding and rating the high-value cells.

Definition 3.5 — Visit performances

We define the **instantaneous visit performance** $q(t)$ as the normalized average value of the cells visited at round t , i.e.,

$$q(t) = \frac{3}{V_{\max_1} + V_{\max_2} + V_{\max_3}} \langle \mathbf{q}(t) | \mathbf{V} \rangle, \quad (3.1)$$

$$= \frac{3}{V_{\max_1} + V_{\max_2} + V_{\max_3}} \sum_c q_c(t) V_c, \quad (3.2)$$

where $(V_{\max_1} + V_{\max_2} + V_{\max_3})/3$ is the normalization constant so that $q(t) = 1$ when every player opens the three cells of highest values: V_{\max_1} , V_{\max_2} , V_{\max_3} .

In the same way, we introduce the **cumulative visit performance** $Q(t)$, as the normalized average value of the cells visited up to round t , i.e.,

$$Q(t) = \frac{3}{V_{\max_1} + V_{\max_2} + V_{\max_3}} \langle \mathbf{Q}(t) | \mathbf{V} \rangle. \quad (3.3)$$

These observables are normalized so that $q(t) = 1$ and $Q(t) = 1$ correspond to the best theoretical performance, i.e., when every individual would visit the three best cells of the grid at round t and up to round t . Note that the cumulative visit performance at the end of the game, $Q(t = N_{\text{rounds}})$, is directly proportional to the normalized score of the group, \hat{S} .

Similarly, one can define the corresponding observables for the ratings instead of the visits.

Definition 3.6 — Rating performances

We define the **instantaneous rating performance** $p(t)$ as the normalized average value of the cells visited by participants weighted by their ratings at round t , i.e.,

$$p(t) = \frac{1}{V_{\max_1}} \langle \mathbf{p}(t) | \mathbf{V} \rangle, \quad (3.4)$$

$$= \frac{1}{V_{\max_1}} \sum_c p_c(t) V_c, \quad (3.5)$$

where V_{\max_1} is a normalization constant so that $p(t) = 1$ when every star is on the highest value cell of value V_{\max_1} .

In the same way, we define the **cumulative rating performance** $P(t)$ as the normalized average value of the cells visited weighted by their ratings up to round t , i.e.,

$$P(t) = \frac{1}{V_{\max_1}} \langle \mathbf{P}(t) | \mathbf{V} \rangle. \quad (3.6)$$

In cases where the vectors $\mathbf{p}(t)$ or $\mathbf{P}(t)$ are undefined, we adopt the convention of setting $p(t) = 0$ and $P(t) = 0$ to be consistent with the fact that there are no stars.

These two quantities measure whether the distribution of stars is concentrated on the high-value cells or not. The closer $p(t)$ and $P(t)$ are to 1, the more stars are placed on high-value cells, and these quantities will be equal to 1 if all the stars are on the cell having the highest value ($V_{\max_1} = 99$, in all our experiments).

Hence, $p(t)$ and $P(t)$ quantify the instantaneous and cumulative 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 $t = N_{\text{rounds}}$) indicates that the participants concentrate the allocation of stars on high-value cells. Conversely, a low value of $P(t)$ indicates that participants allocate stars to low-value cells.

It is important to note that a high value of $P(t = N_{\text{rounds}})$ does not imply a high final individual (S) or group (\hat{S}) score. For example, if all the players play randomly but put only stars in the cell of value higher than 60, they will have a low score but $P(t = N_{\text{rounds}})$ will be larger than $60/99 \approx 0.6$. Conversely, if the players achieved a very high group score \hat{S} (implying that they found and visited very high-value cells), but decided to only put stars in cells of value less than 10, one would observe a final $P(t = N_{\text{rounds}}) \leq 10/99 \approx 0.1$

3.2.4 Inverse participation ratio

Another useful observation would be determining whether players visit and rate the same cells or explore different ones. In other words, we want to understand if the visits and ratings are concentrated on a few specific cells or distributed over the whole grid. To explore this aspect, we introduce the concept of the Inverse Participation Ratio (IPR), which is used in quantum mechanics to measure the localization of quantum states [Kramer and MacKinnon, 1993].

Definition 3.7 — Inverse Participation Ratio (IPR)

Let \mathbf{x} be a probability vector. Then, the inverse participation ratio $\text{IPR}(\mathbf{x})$ of \mathbf{x} is defined as:

$$\text{IPR}(\mathbf{x}) = \frac{1}{\sum_c x_c^2}. \quad (3.7)$$

The IPR serves as a measure to quantify the concentration or spread of a probability distribution over a set of elements. It provides an idea of the effective number of elements on which the probability distribution is concentrated. In a scenario where the distribution is equally spread across n cells out of a total of N , with each cell having a probability of $1/n$, the IPR value becomes $1/[n \times (1/n)^2] = n$. This indicates that the distribution is spread across n cells. In a more general case, the IPR measures the concentration or dispersion of a probability distribution, with lower values indicating more concentration

3.2. Quantifying the Collective Behavior

on fewer elements, while higher values imply a more even spread across multiple elements.

In our case, we are interested in the IPR of the vectors $\mathbf{q}(t)$, $\mathbf{Q}(t)$, $\mathbf{p}(t)$ and $\mathbf{P}(t)$. $\text{IPR}(\mathbf{q}(t))$ and $\text{IPR}(\mathbf{Q}(t))$ indicates the effective number of distinct cells that the players have explored or visited at a specific round t and cumulatively up to round t , respectively. Similarly, the IPR of $\text{IPR}(\mathbf{p}(t))$ and $\text{IPR}(\mathbf{P}(t))$ represents the effective number of cells that contain stars at a particular round t and up to round t , respectively. In the case where no stars have been deposited on the grid (in particular, at $t = 0$), $\mathbf{p}(t)$ or $\mathbf{P}(t)$ are undefined, and we adopt the convention $\text{IPR}(\mathbf{p}(t)) = 0$ and $\text{IPR}(\mathbf{P}(t)) = 0$.

3.2.5 Fidelity

Another valuable information concerns the relationship between the hidden values of the cells in the grid and the fraction of visits or ratings received by these cells up to round t . To quantify this relationship, we introduce the fidelity, denoted F , which is a measure of the relationship between two vectors.

Definition 3.8 — Fidelity of two vectors

Let \mathbf{x} and \mathbf{y} be two vectors. Then, the fidelity $F(\mathbf{x}, \mathbf{y})$ of \mathbf{x} to \mathbf{y} , is defined as

$$F(\mathbf{x}, \mathbf{y}) = \frac{\langle \sqrt{\mathbf{x}} | \sqrt{\mathbf{y}} \rangle}{\|\sqrt{\mathbf{x}}\|_2 \|\sqrt{\mathbf{y}}\|_2}, \quad (3.8)$$

where the vector $\sqrt{\mathbf{x}}$ is defined as follows: For each component x_c in \mathbf{x} , the corresponding component of $\sqrt{\mathbf{x}}$ at index c is given by $\sqrt{x_c}$.

The fidelity F lies in the range $[0, 1]$. A perfect fidelity of 1 is achieved when the two vectors are strictly proportional. A fidelity of 0 is achieved when the vectors are perpendicular. For vectors with positive entries, this corresponds to the support of their non-zero entries being disjoint.

In our experiment, we are interested in the fidelity of the hidden values of the grid to the ratings, $F(\mathbf{P}(t), \mathbf{V})$, and to the visits, $F(\mathbf{Q}(t), \mathbf{V})$. Since $\mathbf{P}(t)$ and $\mathbf{Q}(t)$ are probability vectors, and that for a probability vector \mathbf{x} , we have $\sum_c x_c = 1$, hence $\|\sqrt{\mathbf{x}}\|_2 = \sqrt{\sum_c (\sqrt{x_c})^2} = \sqrt{\sum_c x_c} = 1$. Therefore, the fidelity of a probability vector \mathbf{x} to the value vector \mathbf{V} is

$$F(\mathbf{x}, \mathbf{V}) = \frac{\langle \sqrt{\mathbf{x}} | \sqrt{\mathbf{V}} \rangle}{\|\sqrt{\mathbf{V}}\|_2}, \quad (3.9)$$

$$= \sum_c \sqrt{\frac{x_c V_c}{\sum_{c'} V_{c'}}}, \quad (3.10)$$

where \mathbf{x} represents either $\mathbf{Q}(t)$ or $\mathbf{P}(t)$.

The choice to use the square root of the elements rather than a simple correlation is based on the properties of probability distributions, where using the square root is more natural and common when dealing with their norms. This approach aligns with the Hellinger distance, a well-known measure in probability theory that emphasizes differences

between distributions. Fidelity measures the alignment between two vectors through their scalar product and is closely related to the Hellinger distance H between two probability distributions [Hellinger, 1909; Pollard, 2001]. Indeed, for two probability vectors \mathbf{x} and \mathbf{y} (so that $\|\sqrt{\mathbf{x}}\|_2 = \|\sqrt{\mathbf{y}}\|_2 = 1$), we have

$$H^2(\mathbf{x}, \mathbf{y}) = \frac{1}{2} \|\sqrt{\mathbf{x}} - \sqrt{\mathbf{y}}\|_2^2, \quad (3.11)$$

$$= \frac{1}{2} \sum_c (\sqrt{x_c} - \sqrt{y_c})^2, \quad (3.12)$$

$$= 1 - \sum_c \sqrt{x_c y_c}, \quad (3.13)$$

$$= 1 - F(\mathbf{x}, \mathbf{y}), \quad (3.14)$$

showing that the fidelity is given by $F = 1 - H^2$.

Note that if $\mathbf{P}(t)$ is undefined (in particular, at the very beginning of the game), we adopt the convention $F(\mathbf{P}(t), \mathbf{V}) = 0$ so that the fidelity begins at 0 and increase throughout the game.

In the context of real-world 5-star rating systems, a high fidelity of the cumulative ratings $\mathbf{P}(t)$ would indicate that the ratings provide a fair representation of the actual value of the different options, which is, in principle, the desired outcome for online merchants and services. 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 in detail.

3.3 Quantifying Individual Visit Behaviors

Let us now examine the observables used to quantify the cell-visiting behavior of individuals. These observables are determined at the game level by calculating averages over each player's behavior, focusing only on the cells they visited.

3.3.1 Normalized individual score

In the same way that the normalized group score is defined (see Definition 3.1), it is possible to define a normalized individual score.

Definition 3.9 — Normalized individual score

The normalized individual score of a player, denoted S , is defined as the sum of the values of the cells visited by the player during the game divided by the maximum possible score. This maximum score corresponds to the score a player would achieve if they selected the three best cells in each round, resulting in a total value of $(V_{\max_1} + V_{\max_2} + V_{\max_3}) \times N_{\text{rounds}}$, where V_{\max_1} , V_{\max_2} , and V_{\max_3} are the three cells with the highest value, and N_{rounds} is the number of rounds.

This normalized individual score serves as a valuable metric for assessing the performance of individuals in a game. It provides a simple way of comparing individuals with

3.3. Quantifying Individual Visit Behaviors

each other, enabling a clear evaluation of their relative performance.

In the same way that the mean normalized group score is defined (see Definition 3.2), we can define the mean normalized individual score.

Definition 3.10 — Mean normalized score

The mean normalized score, denoted as $\langle S \rangle$, is defined as the arithmetic average of the normalized individual scores over all games and over all players of these games, for a given experimental condition.

Note that the mean normalized individual score, $\langle S \rangle$, and the mean normalized group score, $\langle \hat{S} \rangle$, are identical, as every group is composed of the same number of players. Therefore, we will refer to the mean normalized score of an individual or group as simply the mean normalized score, and use the notation $\langle S \rangle$.

3.3.2 Values of visited cells

Natural quantities of interest are the mean values of the three cells visited by a player, which leads us to the following definition.

Definition 3.11 — Values of visited cells

Let $V_1(t)$, $V_2(t)$, and $V_3(t)$ denote the mean value of the highest, second-highest, and third-highest cell opened by the players at round t , averaged over all considered games and all players.

These observables provide insights into the dynamic evolution of cell values explored by players throughout a game. Additionally, they highlight the disparities among the values of the three different cells.

It is worth noting that the sum of these three values, $V_1(t) + V_2(t) + V_3(t)$, corresponds, up to a scaling factor, to the normalized average value $q(t)$ of the cells visited during round t (see Definition 3.5).

In the same way that we have defined instantaneous and cumulative observables in the previous section, we can define the average value of the cells visited by a player with the highest values since the beginning of the game.

Definition 3.12 — Values of the cells visited with the highest values since the beginning

Let $V_{B1}(t)$, $V_{B2}(t)$, $V_{B3}(t)$ be respectively defined as the mean value of the highest, second-highest, and third-highest distinct cell opened by a player from the beginning up to round t , averaged over all considered games and all players.

Note that, by definition, we have for all $i = 1, 2, 3$, $V_{Bi}(t) \geq V_i(t)$. The difference in value between these quantities and the previous ones characterizes the tendency of the players to revisit their best cells. A significant difference suggests that players find high-value cells but do not revisit them frequently, whereas a small difference indicates a higher tendency to revisit these cells.

3.3.3 Probability of revisiting cells from the previous round

To more precisely measure players' inclination to revisit cells, we propose introducing novel metrics or observables.

Definition 3.13 — Probability of revisiting cells visited in the previous round

For each round $t > 1$, let $B_1(t)$, $B_2(t)$, and $B_3(t)$ be respectively the average probability to revisit the cell from the previous round ($t - 1$) with the highest, second-highest, and third-highest value.

These probabilities provide valuable insight into the likelihood that players revisit cells from the previous round. In particular, when a notion of score is introduced in the game, we expect that players will revisit more often cells as they discover better and better cells, so that $B_1(t)$, $B_2(t)$, and $B_3(t)$ are expected to growth with the round t .

3.3.4 Probability of finding cells with the highest values

Since the players were instructed to find the high-value cells, it is also important to assess whether they managed to find these cells or not. To quantify this, let us introduce the probability of finding a particular cell with a value V , averaged over all considered games and all their players.

Definition 3.14 — Probability of finding cells with given values

For a given cell c (of value V_c), we define $\mathbb{P}_c(t)$ as the probability that a player has visited that specific cell at least once from the start of the game until round t .

Let \hat{V} be a set of values. Then, for each round t , we define $\mathbb{P}_{\hat{V}}(t)$ as the average of the probability of finding a cell with value in \hat{V} :

$$\mathbb{P}_{\hat{V}}(t) = \frac{\sum_{c|V_c \in \hat{V}} \mathbb{P}_c(t)}{\sum_{c|V_c \in \hat{V}} 1}, \quad (3.15)$$

where the denominator is equal to the number of cells whose value is in the set \hat{V} .

When analyzing our experiments, we will consider the probability $\mathbb{P}_{\hat{V}}(t)$ for the following sets of values: $\hat{V} = \{99\}$ consisting of the best cell; $\hat{V} = \{84, 85, 86\}$ consisting of the four next best cells (one cell with value 84 and 85; two cells with value 86); and $\hat{V} = \{71, 72\}$ also consisting of four cells (two cells with each value). In some cases, we will also consider the average probability of finding these nine cells, associated to $\hat{V} = \{99, 84, 85, 86, 71, 72\}$

We can compare the observed probabilities of finding particular cells with the theoretical probability of finding a specific cell obtained by exploiting two simple strategies. First, consider a strategy where a player opens N_{turns} cells randomly at every round t , with possible random revisits at successive rounds. With this strategy, the probability $\mathbb{P}_{\text{random}}(t)$ of finding a specific cell at least once from the start of the game until round t

3.4. Quantifying Individual Rating Behaviors

is given by

$$\mathbb{P}_{\text{random}}(t) = 1 - \Pr(\text{Not finding the cell up to round } t), \quad (3.16)$$

$$= 1 - (\Pr(\text{Not finding the cell at round } t))^t, \quad (3.17)$$

$$= 1 - \left(\prod_{i=0}^{N_{\text{turns}}-1} \frac{N-i-1}{N-i} \right)^t, \quad (3.18)$$

$$= 1 - \left(\frac{N - N_{\text{turns}}}{N} \right)^t. \quad (3.19)$$

Next, consider a better strategy where the player sequentially opens all the cells on the grid, without repetition. Since the specific cell to find in the grid is unknown to the player, the player could start by opening the cell in the upper-left corner, then the next cell on the right, and after finishing a row, proceed to the next row. Actually, if the player were playing the game alone, this strategy would be optimal for finding a particular cell (see also Figure 5.3 for a nice illustration of the biased randomness in humans). In fact, in our experiments, we have observed some real players using this strategy in the first few rounds of a game, especially when playing alone on a grid ($N_{\text{players}} = 1$; instead of $N_{\text{players}} = 5$, as in most of our experimental conditions). With this strategy, the probability $\mathbb{P}_{\text{sequential}}(t)$ of finding a specific cell at least once from the start of the game until round t is given by

$$\mathbb{P}_{\text{sequential}}(t) = 1 - \prod_{i=0}^{N_{\text{turns}}t-1} \frac{N-1-i}{N-i} = \frac{N_{\text{turns}}}{N} \times t. \quad (3.20)$$

This formula holds as long as $N_{\text{turns}}t \leq N$, while for $N_{\text{turns}}t > N$, the player has already visited all available cells, and $\mathbb{P}_{\text{sequential}}(t) = 1$.

When we address the different experimental conditions studied in this thesis, we will compare the observed $\mathbb{P}_{\hat{v}}(t)$, to the prediction $\mathbb{P}_{\text{random}}(t)$ and $\mathbb{P}_{\text{sequential}}(t)$ for these two simple strategies. This comparison will complete our analysis of the degree of collaboration and deception in the game, as well as the tendency of players to replay their best cells. Indeed, collaboration is expected to boost the probability $\mathbb{P}_c(t)$ of finding a given high-value cell c , while revisiting cells limits the exploration of the grid by the players and should lower this probability.

3.4 Quantifying Individual Rating Behaviors

Now, let us focus on the second part of the player's turn: the rating of the cells. The observables described here are computed using the values of the cells opened and their ratings, i.e., the number of stars used by the player to rate cells. In this part, we are not studying the time dependence but rather, the average behavior of the players in the game.

3.4.1 Mean rating

The simplest observable that can be used to study the rating of the cells is the mean rating as a function of the value of the cell.

Definition 3.15 — Mean rating

The mean rating $R(V)$ for a player or a group of players of a cell with a value of V , is defined as the arithmetic average of the ratings given by that player or that group to all cells with the value V . If there are multiple cells with the same value V in the grid, the ratings of all such cells are included in the calculation of the average.

This is the most fundamental quantity to study the rating behavior, as it allows us to characterize and understand the rating behavior of a player or of a (sub)group of players.

Note that in our experiments, each player is limited to visiting a maximum of $N_{\text{rounds}} \times N_{\text{turns}} = 60$ different cells during a game. Consequently, players rarely encounter every unique value during the game. To deal with this limitation, I implemented a binning process where close values are grouped. See Table 3.1 for the specific binning values. This approach also serves to reduce the error bars associated with the limited set of values encountered.

0, 1, 2, 3	1.59
4, 5, 6, 7, 8, 9	5.91
11, 12, 13, 14	12.75
19, 20, 21, 22, 24	21.08
27, 28	27.25
43, 44, 45, 46	44.25
51, 53	52.50
71, 72	71.50
84, 85, 86	85.25
99	99.00

Table 3.1: **Binning table.** The table shows a mapping of values, where the left column represents the grouped values and the corresponding right column represents their new common value. The new value is determined by calculating the weighted average of all grouped values, taking into account their respective frequencies within the grid. For example, for the first row of the table, the game has 24 cells with a value of 0, 44 with a value of 1, 47 with a value of 2, and 32 with a value of 3. This results in a new value of $(24 \times 0 + 44 \times 1 + 47 \times 2 + 32 \times 3)/(24 + 44 + 47 + 32) = 1.59$.

3.4.2 Fitting of the mean rating

The mean rating of cells as a function of their values is a very important quantity since it can reveal different behavioral strategies. To facilitate the analysis of these distinct behaviors, we will examine a linear approximation of the mean rating for each player. This simplification condenses the complexity into two numerical values.

3.5. Computation of the Error Bars: the Bootstrap Method

Definition 3.16 — Intercept and slope parameters of the mean rating

Consider the following linear approximation function:

$$\widehat{R}(V) = \hat{u}_0 + 5 \hat{u}_1 \frac{V}{V_{\max_1}}, \quad (3.21)$$

where V_{\max_1} represents the highest value in the grid, and \hat{u}_0 and \hat{u}_1 denote two parameters.

Subsequently, u_0 and u_1 correspond to the optimized values of the parameters \hat{u}_0 and \hat{u}_1 obtained by applying the least squares method to fit the above linear function with the mean rating. u_0 is referred to as the intercept parameter, while u_1 is referred to as the slope parameter for the mean rating.

In our experiments, the typical range of u_0 lies between -1 and 6 , while for u_1 , the range is between -2 and 2 .

While these two parameters only provide a linear approximation of the mean rating, they manage to capture a wide range of behavior. For example, an individual with $u_0 = 0$ and $u_1 = 1$ rates a cell of value 0 with an average of 0 star, and rates the cell with the maximum value $V_{\max_1} = 99$ with an average of 5 stars. On the other hand, another individual with $u_0 = 5$ and $u_1 = -1$ assigns 5 stars to cells with value 0 and 5 stars to the cell with value V_{\max_1} . In addition, an individual with $u_0 = 0$ and $u_1 = 0$ simply does not allocate any star to any cell, and in a sense, refuses to participate in the rating process.

3.4.3 Fractions of ratings

The mean rating provides the average number of stars assigned to cells with specific values. However, it does not provide information about the distribution of each rating. To conduct a more comprehensive analysis, we have to consider the fractions of ratings, as different rating distributions can produce the same mean rating.

Definition 3.17 — Fractions of ratings

The fractions of ratings, denoted as $P_s(V)$, are defined as the proportion of instances where cells with value V received an s -star rating, relative to all instances of cells with value V being rated. These fractions can be interpreted as the probability for a player to assign an s -star rating to a cell of value V .

As with the mean value, this quantity can be grouped into bins to reduce observational errors and account for the possibility that not all values may be observed during a game.

3.5 Computation of the Error Bars: the Bootstrap Method

In the numerous graphs presented in this thesis, there are two types of data: experimental data and simulated data obtained from various models. Each observable associated with the experimental data will be shown with error bars corresponding to a confidence level of 68.27% , which are evaluated using a bootstrap method (detailed below). The significance of the 68.27% confidence level lies in its relationship to the standard deviation of a normal

distribution, where approximately 68.27% of the data falls within one standard deviation from the mean. On the other hand, the simulation results are presented without error bars since they correspond to the average over 1,000,000 simulated games, resulting in a negligible error relative to the scales of the graphs.

When it comes to estimating the statistical properties of an observable using a sample of size N drawn from a population, several methods come into play. For instance, when calculating the standard error of the mean for an observable, the central limit theorem provides an exact estimate for large enough N : σ/\sqrt{N} , where σ is the sample standard deviation, which is assumed to be finite. However, we will deal with data of moderate size, and with several observables with skewed distributions and/or strictly constrained to belong to a specific interval. In fact, this is the case of all observables measuring a probability, and other observables like the fidelity, F , which is also constrained between 0 and 1. For instance, if some average probability is found to be small, it is expected that its lower confidence interval is smaller than its upper confidence interval, due to the strict lower bound at 0 for a probability. The bootstrap method proves to be a robust solution for dealing with such data, alleviating the need to assume normally distributed data or rely on the law of large numbers.

The bootstrap method, or simply “bootstrapping”, is a resampling technique designed to evaluate some characteristics, such as confidence intervals, of statistical parameters, such as mean or median, derived from an unknown probability distribution. The method consists in iteratively generating new samples by drawing with replacement from an original sample [Efron, 1979; Davison and Hinkley, 1997].

In practice, getting the confidence interval of some statistical parameter, such as the mean or the median, of a sample of N observables using the bootstrap method consists of the following steps, which are illustrated in Figure 3.1:

1. Start with some dataset of size N for some observed quantity.
2. Iteratively draw with replacement M new samples, also of size N , from the original dataset. These artificial samples are called “bootstrap samples”. Note that some elements of the original dataset may appear multiple times within a bootstrap sample, while others may not be included at all. For robust results, a sufficient number of bootstrap samples M is required. In the context of this thesis, I used 10,000 bootstrap samples.
3. For each of these bootstrap samples, compute some statistical parameters, which are called “bootstrap estimate”.
4. Construct the distribution of these M bootstrap estimates.
5. Use this distribution to derive the confidence intervals of the considered statistical parameters for the observed quantity.

Therefore, this method allows the estimation of sampling distribution for various statistics, even in scenarios involving unknown probability distributions and limited data size.

3.5. Computation of the Error Bars: the Bootstrap Method

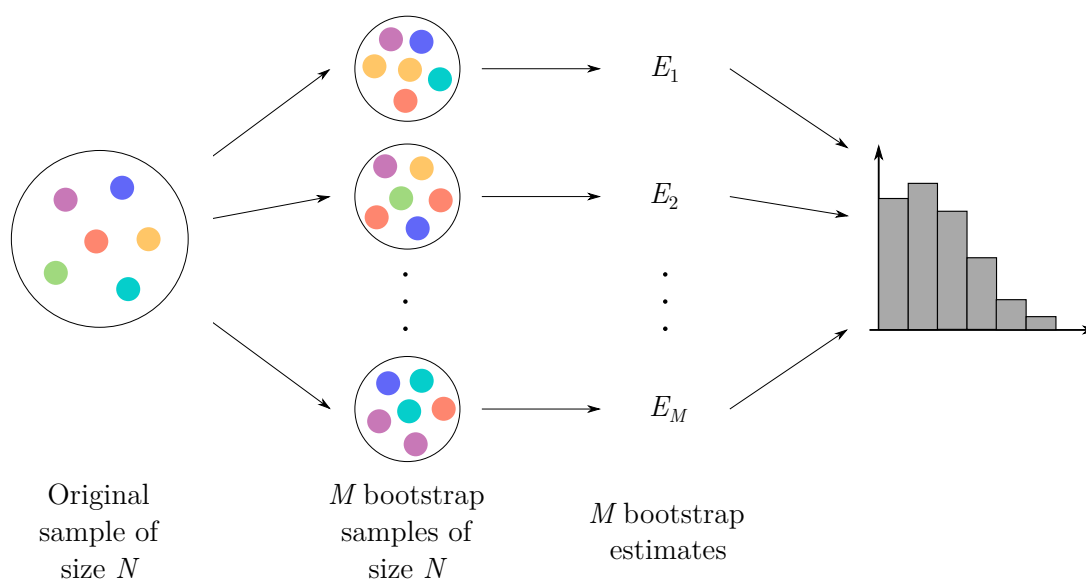


Figure 3.1: **Illustration of the bootstrap method.** From an original sample of size N is drawn with replacement M new samples, called bootstrap samples, of size N . For each of these bootstrap sample, some statistical parameter, called bootstrap estimate, is computed. Then the distribution of these bootstrap estimates is used to derive estimates for the confidence interval of the statistics.

Chapter Summary

This chapter presents the stochastic agent-based model used to simulate player behavior in the STIGMER game. The model consists of agents with different strategies that govern their actions during the game. In particular, each agent's strategy is decomposed into two independent components: the visit strategy, which governs cell selection, and the rating strategy, which governs cell rating.

During a game, players have access to both collective and private information. The collective information is represented by the color of the cells on the grid, i.e., the ratings, while the private information corresponds to the player's knowledge about their previously visited cells.

The visit strategy considers both types of information available to agents. This allows them to either revisit cells that were already visited in the previous round, using private information, or to visit other cells based on their color, using collective information. When using collective information, agents can choose to focus on either dark cells, following social information, or on white cells, for exploration of unrated cells.

The rating strategy assigns a probability to give a certain star rating to a visited cell, which generally depends on the value of the cell.

The model is parameterized, allowing for the generation of agents with different behaviors, by adjusting its parameter values. These variations include agents that explore extensively, those that avoid exploration, and agents that engage in collaborative or deceptive rating practices.

The model is then applied to create different types of agents. "Mimic agents" that mimic human behavior, "Optimized agents" that are tailored for specific tasks and environments, and agents with specific strategies that are used to understand the impact of these strategies on individual and collective performance.

All parameters governing these agents are either specified manually, or determined by fitting the characteristics of human behavior, or determined by exploiting optimization methods (in particular, Monte Carlo).

Contents

4.1	General Description of the Model	52
4.2	Visit Strategy	53
4.2.1	Revisiting the cells visited during the previous round	53
4.2.2	Exploration of the grid	54
4.3	Rating Strategy	55
4.3.1	Using the mean rating	55
4.3.2	Direct specification of the distribution functions	55
4.4	Types of Agents	57
4.4.1	Mimic agents	57
4.4.2	Optimized agents	57
4.4.3	Other agents	57
4.5	Determination of Model Parameters	57
4.5.1	Fitting	58
4.5.2	Monte Carlo method	58

This chapter introduces a stochastic agent-based model of the STIGMER game that will then be used in the following chapters. First, I present the general idea of the model, which is divided into two parts: the visit strategy and the rating strategy. Then, I introduce the different types of agents, and finally, I talk about the determination of the parameters of the model.

The agents developed in this model will be used in the rest of the thesis for several purposes. They will allow us to quantitatively identify the strategies of our human subjects for visiting and rating cells, and thereby understand their distinct effects both on individual and collective performance. The model will also allow us to find the best strategies in different situations. Finally, we will address experiments where a human subject is playing against four virtual agents (bots), allowing for full control of each participant's experimental environment.

4.1 General Description of the Model

To better understand the impact of different environmental situations on the way humans exploit the traces in our experiment, we have developed a stochastic agent-based model.

The model follows the actual game rules. That is, for each simulated game, a group of N_{players} independent agents plays the game for N_{rounds} rounds, in which they must visit and evaluate N_{turns} cells. The agents themselves consist of two distinct and independent parts: *(i)* the agents' strategy for visiting cells, i.e., which cells an agent decides to visit in each round; *(ii)* their strategy for evaluating the visited cells, i.e., how many stars an agent should assign to a cell of a given value.

4.2. Visit Strategy

The visit strategy and the rating strategy, which will be presented in the next sections, are designed to be simple while preserving the most important features of real players' strategies. These strategies will depend on several parameters and are therefore not strategies in the sense of game theory, but rather two classes of strategies that allow for an infinite number of strategies by tuning the value of the parameters.

4.2 Visit Strategy

When deciding which cell to visit, players have two options. They can either use their personal information, which consists of their knowledge of previously visited and rated cells, or they can use the collective information embedded in the grid colors, which reflect the ratings left by all players since the beginning of a game.

The strategy used by the agents to visit cells is organized around two main options. Firstly, agents have a probability of revisiting the cells visited during the previous round using their private information. Secondly, if they do not revisit cells, they will explore other cells in the grid using the collective information. Note that in the first round, as there is no public or private information available, the cells are selected at random.

This strategy relies on the cumulative fraction of stars, represented by the colors, and the values of the cells visited during the previous round. The equations do not explicitly include a round-specific dependency. However, there is an implicit temporal dependency because the cumulative fraction of stars and the values of the cells visited at the previous round change over time.

4.2.1 Revisiting the cells visited during the previous round

Except for the first round ($t > 1$), agents adopt the following strategy. For each of the three cells $i = 1, 2, 3$ to visit, they either visit the i -th cell with the highest value 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 $1 - P_i^R(V_i(t-1))$.

For each cell $i = 1, 2, 3$, this probability is given by:

$$P_i^R(V_i(t-1)) = \begin{cases} 0 & \text{if } V_i(t-1) < a_i \\ \frac{V_i(t-1) - a_i}{V_{\max_1}} b_i & \text{if } a_i \leq V_i(t-1) < a_i + \frac{V_{\max_1}}{b_i} \\ 1 & \text{otherwise} \end{cases}, \quad (4.1)$$

where a_i and $b_i > 0$ are parameters, and $V_{\max_1} = 99$ is the maximum cell value on the grid. Refer to Figure 4.1 for a graphical representation of the curve. Therefore, an agent never revisits a cell of value $V_i(t-1) < a_i$ and always revisits a cell of value $V_i(t-1) > a_i + V_{\max_1}/b_i$. Between these two thresholds, the probability of revisiting the i -th cell with the highest value linearly interpolates between 0 and 1.

Each of these three probabilities depends on two parameters, for a total of six parameters. By adjusting the values of these parameters, it is possible to reproduce a wide range of behaviors. For example, when $a = 0$ and $b = 1$, the resulting probability is strictly proportional to the value of the cell and saturates at 1 exactly at $V = V_{\max_1}$. Moreover,




Figure 4.1: **Probability to revisit a cell as a function of its value.** The probability $P^R(V)$ to revisit a cell as a function of its value V , highlighting the influence of the parameters a and b .

if $a = V_{\max 1}$, a cell is never revisited. Finally, choosing $a = V_{\max 1}/2$ and $b \gg 1$ ensures that revisits never occur below the value $V_{\max 1}/2$, but always beyond that value.

Note that we have assumed that the agents have only a one-round memory of their visited cells. This hypothesis is supported by the data analysis, which found that if individuals did not choose to revisit a particular cell in one round, they had a very small chance of revisiting it in subsequent rounds.

4.2.2 Exploration of the grid

When agents do not visit one of the cells visited in the previous round, they explore other cells in the grid. This is done by associating to each cell c a probability $P^E(c, t)$ to be selected at round t :

$$P^E(c, t) = \varepsilon \frac{1}{N} + (1 - \varepsilon) \frac{P_c^\alpha(t-1)}{\sum_{c'} P_{c'}^\alpha(t-1)}, \quad (4.2)$$

where $P_c(t-1)$ is the cumulative fraction of stars deposited in cell c up to time $t-1$ (see Definition 3.4), and $\varepsilon \in]0, 1]$ and $\alpha > 0$ are two parameters.

The defining equation for $P^E(c, t)$ presents two different contributions, each with a clear interpretation. The first contribution, ε/N , translates the fact that an agent has a probability ε of picking a cell completely at random (uniform random exploration). The second contribution, $(1 - \varepsilon)P_c^\alpha(t-1)/\sum_{c'} P_{c'}^\alpha(t-1)$, corresponds to the weight of the social information embedded in the cell colors, which are increasing functions of P_c (see Eq. (2.1)). This weight is determined by the fraction of stars in each cell: the higher the fraction, the higher the weight. The exponent α controls the contrast of the selection of a colored cell on the grid. A high value for α would correspond to a strong preferential selection of the highly marked cells, while a small value for α would lead to a more homogeneous selection of cells among the marked ones.

If there are no stars on the grid, the cumulative fraction of stars will be undefined. In this case, a cell will be selected randomly. Moreover, to avoid cells being selected twice, if the selected cell is one of the cells visited in the previous round or if the cell has already been visited during this round, another one is randomly selected by the same procedure.

4.3 Rating Strategy

The second step of the model addresses the rating strategy, i.e., the number of stars an agent uses to rate a cell as a function of its value.

For simplicity, we have made two hypotheses. The first one is that the rating strategy does not depend on the round, nor the number of cells already opened in the round. The second one is that the rating strategy does not depend on the cumulative fraction of stars of the cell, i.e. its color. Therefore, the rating strategy is only a function of the value of the cell visited.

The rating strategy is inherently straightforward: ratings are established through a random process governed by a discrete probability distribution which depends on the value of the cell. This distribution, denoted as $P_s(V)$, assigns probabilities to rate a cell of value V with s stars ($s = 0, 1, \dots, 5$). The determination of these probabilities can be done using the two approaches described below. Each method has some parameters that allow for a large variety of behaviors.

4.3.1 Using the mean rating

The first approach to determining the probability distributions is to use the mean rating. In this method, if the mean rating assigned to a value V is denoted as $R(V)$, then the probability $P_s(V)$ to rate s stars the cell is defined by the following expression:

$$P_s(V) = \begin{cases} 1 - \{R(V)\} & \text{if } s = \lfloor R(V) \rfloor \\ \{R(V)\} & \text{if } s = \lfloor R(V) \rfloor + 1 \\ 0 & \text{otherwise} \end{cases}, \quad (4.3)$$

where $\lfloor \cdot \rfloor$ is the integer part and $\{\cdot\}$ is the fractional part.

Using this method, let us consider the following scenario where an agent has to rate a cell of value 20 with a mean rating, $R(20) = 3.4$. In this case, $\lfloor R(V) \rfloor = 3$ and $\{R(V)\} = 0.4$, resulting in the following probabilities: $P_3(20) = 0.6$, $P_4(20) = 0.4$, and $P_s(20) = 0$ for $s = 0, 1, 2, 5$.

This approach will be particularly useful for constant or linear mean ratings. However, one limitation of this approach is that only two ratings (or one for integer values) have non-zero probabilities for a given mean rating. This results in a lack of diversity in ratings, which contrasts with the diverse rating patterns observed in experiments.

4.3.2 Direct specification of the distribution functions

The alternative method, which will be the most used one, addresses the issue of rating diversity by directly specifying the functions that govern the probability distributions $P_s(V)$. These functions can be any function that takes a value between 0 and 1. In our case, we will be using three different types of functions: sigmoid-like functions, Gaussian-like functions, and linear functions.

The probability $P_s(V)$ for a cell of value V to be rated with s stars is given by the following equations:




Figure 4.2: **Functions characterizing the rating probability distributions.** Illustration of the role of the parameters in the functions characterizing the probability distributions of rating a cell in the case of (A) the sigmoid-like function, (B) the Gaussian-like function, and (C) the linear function.

- Sigmoid-like functions have the following form:

$$P_s(V) = c_s + d_s \tanh\left(\frac{v - e_s}{V_{\max_1}} f_s\right), \quad (4.4)$$

where $c_s, d_s > 0$, e_s , and f_s are parameters. The parameters e_s and c_s are the x and y coordinates of the inflection point of the sigmoid, respectively. Meanwhile, d_s is the amplitude of the sigmoid, and f_s controls the steepness of the sigmoid curve and can be either positive or negative. See Figure 4.2A.

- Gaussian-like functions are defined by the expression:

$$P_s(V) = d'_s \exp\left(-\left(\frac{v - e'_s}{V_{\max_1}} f'_s\right)^2\right), \quad (4.5)$$

where $d'_s > 0$, e'_s , and $f'_s > 0$ are parameters. The parameter d'_s is the amplitude of the peak, the parameter e'_s is the position of the peak, and finally the parameter f'_s controls its standard deviation. See Figure 4.2B.

- Finally, the linear functions have the form:

$$P_s(V) = c_s'' + f_s'' \frac{V}{V_{\max_1}}, \quad (4.6)$$

where c_s'' and f_s'' are parameters. The parameter c_s'' is the y -intercept and the parameter f_s'' controls the slope which can be positive or negative. See Figure 4.2C.

When using this method, it is crucial to ensure that the sum of probabilities for all values V equals one: $\forall V, \sum_s P_s(V) = 1$. As a result, some of the parameters are interdependent and often require defining the probability of placing one star as the complementary probability of the others. For example, $P_5(V) = 1 - P_0(V) - P_1(V) - P_2(V) - P_3(V) - P_4(V)$. During this process, it is essential to check that the resulting probability, $P_5(V)$ in the example, remains non-negative.

4.4 Types of Agents

In this thesis, the model will be used to create different types of agents. The first type, called “Mimic agents”, are agents that mimic human behavior in a given situation. The second type, called “Optimized agents”, are agents optimized for a specific task in a specific situation. Finally, the last type of agents are not optimized but have a specific strategy

4.4.1 Mimic agents

Mimic agents are agents that mimic human behavior. By mimicking human behavior, these agents serve two purposes. First, they demonstrate that the model introduced above can quantitatively replicate the results observed in our experiments. Second, they help us understand human behaviors, especially their visit and rating strategies, as well as how humans interact with each other through the digital traces they leave on the grid.

For a given experiment, the Mimic agents can be composed of several subtypes of agents with different rating strategies to account for the diversity of behaviors of human individuals. In this case, when a Mimic agent is used in a simulation, one of these subtypes is randomly selected according to some predefined proportions.

4.4.2 Optimized agents

Optimized agents are agents that are optimized for a specific task within a specific environment. These tasks can include goals such as maximizing individual or group scores, or maximizing fidelity. The environment in which the agents are optimized can also vary. Occasionally, agents are optimized while playing with agents that mimic human behaviors. Other times, agents are part of a cohesive group of identical agents. In this case, we optimize both the agent and the group.

While some instances of these agents are limited to theoretical scenarios that focus on optimal strategies for various tasks, other instances are compared to that of Mimic agents. Such comparative analyses provide insights into areas where human decision-making could be improved.

4.4.3 Other agents

The last type of agents are agents that are not optimized but have a specific set of parameters that are interesting to study. For example, we can take our model with reasonable visit strategy parameters and compare several agents with different mean rating slopes. This would allow us to study the effect of the mean rating slope on individual and collective performance.

4.5 Determination of Model Parameters

For the third type of agents that are not optimized, the parameters were put by hand, however, for the Mimic agents and Optimized agents the parameters need to be deter-

mined. To achieve that, I used two main methods: the first one is fitting and the second is a stochastic optimization using a Monte Carlo method.

4.5.1 Fitting

The fitting was only done for the rating strategy of the Mimic. In Section 4.3.2, I introduce three functions that can be directly fitted to the probability distributions of the experimental players. That way, Mimic agents have similar rating strategies to humans.

4.5.2 Monte Carlo method

Another approach to determine the values of the parameters is to use a zero-temperature Monte Carlo method. The Monte Carlo method, which serves as a minimization technique, involves a sequential series of steps. In each step, a small random change is made to a randomly selected parameter. If this change results in a reduction of the quantity to be minimized, the adjusted parameter value is accepted. Conversely, if such a change does not reduce the error, the previous parameter value is kept. The minimization process ends when the error is no longer reduced. To account for the potential existence of several local minima within the error landscape, one must start the Monte Carlo simulations from multiple initial parameter values. The final parameters associated with the best simulation are then retained.

For the Mimic agents, this method was used to determine the optimal parameter values for their visit strategy by minimizing an error defined as the discrepancy between a collection of round-dependent observables, denoted $O_1(t), \dots, O_n(t)$, obtained using the experimental data, and a corresponding set of observables, $\hat{O}_1(t), \dots, \hat{O}_n(t)$, obtained through extensive model simulations (averaging over a substantial 1,000,000 numerical experiments). The error is defined as follows:

$$\Delta = \frac{1}{n} \sum_{i=1}^n \frac{\sum_{t=1}^{N_{\text{rounds}}} (\hat{O}_i(t) - O_i(t))^2}{\sum_{t=1}^{N_{\text{rounds}}} O_i^2(t)}. \quad (4.7)$$

Finally, the derivation of the parameters governing the visit and rating strategies for the Optimized agents involves the application of a similar zero-temperature Monte Carlo procedure, similar to the one described above. However, instead of minimizing an error as previously mentioned, the goal now shifts to either minimizing or maximizing the designated task assigned to these agents. Note that maximizing a quantity can be thought of as minimizing the opposite of that quantity, so the same Monte Carlo method can be used.

Impact of Competitive and Non-Competitive Conditions on Stigmergic Cooperation

Chapter Summary

This chapter presents results exploiting the analysis of [Bassanetti et al. \[2023\]](#), which demonstrate the impact of competition on cooperation and deception. Two types of experiments exploiting the STIGMER game were compared to achieve this. In the first experiment, players were instructed to find high-value cells without any incentives, allowing them to act freely. In the second experiment, a score was calculated as the sum of the visited values. The ranking of the players within the group was determined by their score, and compensation was provided accordingly. Therefore, this second experiment incited players to open high-value cells to achieve a higher score than others, creating an intragroup competition absent in the first experiment.

The analysis indicates that, in the absence of competition, participants explore the grid and find high-value cells. However, in the competitive experiment, which incites participants to visit high-value cells, they revisit previously found high-value cells instead of exploring. This results in a smaller exploration of the grid, leading to participants finding high-value cells less frequently. However, the average value of the cells visited during the game is much higher in this competitive situation.

When examining how individuals rate cells, three categories of behavior can be identified based on their degree of cooperation: collaborators rate proportionally to the cell value; defectors rate inversely proportional to the values; and neutrals give on average the same rating regardless of the value of the cell. In the absence of competition, most individuals tend to behave as collaborators, implying that humans' default behavior is to cooperate in this context. However, the introduction of competition results in far fewer collaborators and in a sharp increase in the fraction of defectors and neutrals. As a result, many individuals tend to withhold their private information and may even spread misinformation.

The model presented in Chapter 4 was used to simulate participant behavior in both experiments. This allowed for a quantitative understanding of human visit and rating strategies. Furthermore, the model allowed to investigate the dynamics of the STIGMER game under various conditions, such as the number of agents, their behavioral strategies, grid size, and number of rounds. It also enabled the investigation of optimal agent strategies based on specific objectives.

Contents

5.1 Introduction	60
5.2 Experimental Design	61
5.3 Collective Behavior	62
5.4 Individual Behavior	66
5.4.1 Visiting cells	67
5.4.2 Rating cells	69
5.4.3 Behavioral profiles of individuals	71
5.5 Model Predictions	77
5.5.1 Modeling human behavior	78
5.5.2 Comparison between theoretical and experimental results	78
5.5.3 Impact of the number of rounds and group size on individual performance and collective dynamics	81
5.5.4 Impact of the rating strategy on agents' performance and the fidelity of ratings	84
5.5.5 Optimization of agents' performance according to specific objectives	86
5.6 Solo Games	92
5.7 Discussion	92

5.1 Introduction

This chapter presents results along the lines of those published in [Bassanetti et al. \[2023\]](#), but for a slightly different experimental setup in its payment structure. However, most conclusions drawn in [Bassanetti et al. \[2023\]](#) will still be valid here. The actual experimental results of [Bassanetti et al. \[2023\]](#) will be considered in Chapter 7 along with several other experimental conditions.

The present chapter first aims to present and understand how individuals use indirect (stigmergic) interactions to interact in the context of our information search task experiment described in Chapter 2. Secondly, this chapter explores how a competitive or non-competitive context impacts cooperation and deception between individuals and influences how they exchange and use the social information, embedded in the traces of their past actions, to perform the information search task.

To begin this exploration, I will first provide a comprehensive overview of the two experiments conducted. I will then delve into an in-depth analysis of both the collective and individual behavior of the players, using the observables introduced in Chapter 3. This analysis is an important step in understanding the behavioral and cognitive mechanisms that drive stigmergic interactions and their impact on collective outcomes. It will also lay the groundwork for understanding how different individuals contribute to the collective task based on their different behavioral profiles, thus differentiating themselves in terms

5.2. Experimental Design

of their level of cooperation and performance. This analysis will also validate the form of the visit and rating strategies implemented in the model introduced in Chapter 4. This model is then exploited to investigate and better understand the mechanisms at work in the experiments. It will allow us to simulate agents that not only mimic human behavior, but also optimize certain quantities, providing valuable insights into the mechanisms at play in both experiments.

5.2 Experimental Design

As stated in the introduction, this chapter presents and compares the results of two experiments that use different reward schemes to explore the impact of competition on cooperation, as well as individual visit and rating strategies.

In both experiments, the experimental sessions involved ten individuals. Each individual started by playing two games of twenty rounds alone. The main goal was to ensure that the participants could familiarize themselves with the web interface and integrate the rules (see below), and to measure their spontaneous behavior when the only information available was the digital trace resulting from their own activity. Participants were then randomly divided into two independent groups of five players and typically played ten games of twenty rounds in the same groups of five players.

In both experiments, the goal of the games was the same: to find the cells of high value. However, there is a difference in the scoring system and the final payout, which leads one experiment to have competition between the participants and the other to have no competition.

In the non-competitive experiment, the scoring system is disabled, and all participants are paid the same amount of 10€. In this context, the actions of the participants (visits and ratings) have no impact on their payout. Therefore, this rule ensures that there is no competition between players.

In the competitive experiment, the score of the participants corresponds to the sum of the values of the cells they open. At the end of an experimental session, the scores obtained by the participant while playing alone and in groups of five were combined to get their total score. Then, in each group, the participants were ranked based on their respective total score, and were paid according to their rank. The first ranked player received 20€, the second received 15€, and the three remaining players (ranked 3–5) received 10€. Therefore, participants are incentivized to have the highest score, which creates competition within the five members of the groups.

However, the two groups of five players are truly independent, and the payment structure does not introduce any form of competition between them. Hence, the fact that there are two independent groups can be seen as a mere way to accumulate twice more data in the same experimental period. This is at variance with the actual payment structure used in [Bassanetti et al. \[2023\]](#) where the ten players of the two groups in the competitive rule were ultimately ranked *together* and paid accordingly. The first ranked player received 20€, the second and third received 15€, and the seven other players (ranked 4–10) received 10€ each. Clearly, this payment structure introduces a mix of intragroup and intergroup




Figure 5.1: **Probability distribution functions of the normalized scores.** Probability distribution functions (PDF) of (A) the normalized individual score S and (B) the normalized group score \hat{S} . The non-competitive Rule 1 is in blue and the competitive Rule 2 is in orange. The dots are the experimental data, and the solid lines are the predictions of the model. The dashed vertical lines represent the mean normalized score $\langle S \rangle$.

competition. This payment structure and others, including a payment structure creating a pure competition between the two groups, will be addressed in Chapter 7. In the rest of the present chapter, we follow the analysis of Bassanetti et al. [2023], but applied to a payment structure where the two groups are completely independent and where there is only intragroup competition.

In the non-competitive experiment, we ran a total of ten sessions. Since each session consisted of two groups of five, this resulted in a total of twenty independent groups, for a total of 200 solo games and 190 five-player games. Similarly, for the competitive experiment, we ran seven sessions, resulting in fourteen independent groups, for a total of 140 solo games and 138 five-player games.

For the rest of this chapter, and following Bassanetti et al. [2023], the non-competitive experiment will be referred to as “Rule 1” and its results will be shown in blue in the figures, while the competitive experiment will be referred to as “Rule 2” and will be shown in orange.

5.3 Collective Behavior

In this section, I analyze the collective performance of groups of individuals using the observable introduced in Chapter 3.

Normalized score

First, let us focus on the normalized scores (see Definitions 3.1 and 3.9). As a reminder, in the non-competitive version of the experiment (Rule 1) there is no score. Therefore, to compare individual and collective performance in both rules, we introduce the normalized score, which is identical in both versions of the experiment, although our human subjects

5.3. Collective Behavior

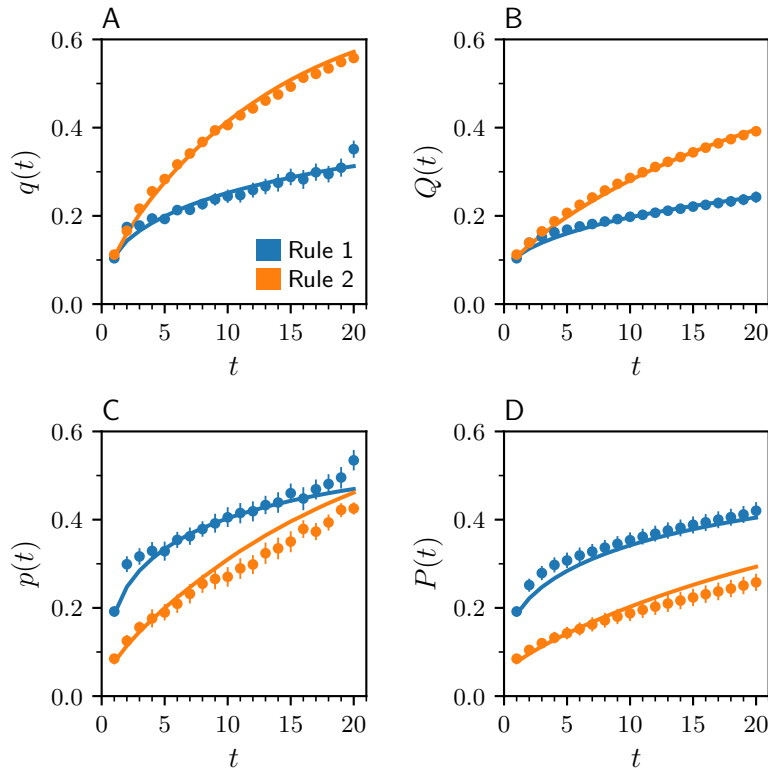


Figure 5.2: **Visit and rating performance.** Average value of the cells visited (A) at round t $q(t)$ and (B) up to round t $Q(t)$. Average value of the cells visited weighted by their ratings (C) at round t $p(t)$ and (D) up to round t $P(t)$. The non-competitive Rule 1 is in blue and the competitive Rule 2 is in orange. The dots are the experimental data, and the solid lines are the predictions of the model.

were only aware of this notion of score while playing Rule 2.

Figure 5.1 shows the probability distribution function of the two normalized scores (individual and group). One can see that the mean normalized score is higher in Rule 2: $\langle S \rangle = 0.40$ vs. $\langle S \rangle = 0.24$. This indicates that the competitive condition provides a stronger incentive to visit the high-value cells.

For reference, as the average value of a cell in the grid is of 9.7, on average individuals playing randomly would get a normalized score of 0.11. Therefore, even in the non-competitive condition, individuals perform twice better than if they were playing randomly.

Visit and rating performance

The analysis of instantaneous and cumulative visit performance (see Definition 3.5) is presented in Figures 5.2A and 5.2B. In both rules, the initial round shows the same average value of the visited cell ($q(0) = Q(0) = 0.10$), since participants visit cells randomly, gaining no information from the blank grid. In fact, Figure 5.3 reveals that the distribution of the visited cells in the first round is not uniform. In subsequent rounds, the average value of $q(t)$ gradually increases as participants explore, visit, and rate cells with higher




Figure 5.3: **Human randomness bias.** Fraction of visits of each cell in the grid during the first round of the experiment. This heatmap combines the distribution of the three cells visited in the first round in all experiments, even some that are not presented in this thesis. We note that the distribution of the cells visited in the first round is not uniform, with a higher density in the corners (especially the upper-left corner), in the center, and along the diagonals. A similar phenomenon was observed in [Fraisse \[1968\]](#). This bias does not affect our experiment, since the value of cells are randomly distributed on the grid, and are shuffled before every game.

values. However, this increase is not uniform in the two rules. Specifically, in Rule 2, individuals tend to favor cells with higher values compared to Rule 1. For example, during the final round, we observe $q(20) = 0.24$ for Rule 1, while Rule 2 has a higher $q(20) = 0.56$. This difference in performance between the two rules is consistent with the higher mean normalized scores in Rule 2 (see Figure 5.1).

While Rule 2 outperforms Rule 1 in terms of visit performance, an interesting contrast emerges when examining the instantaneous and cumulative rating performance (see Definition 3.6 and Figures 5.2C and 5.2D). In fact, the average value of a visited cell weighted by its corresponding rating is higher in Rule 1. For example, in the final round of Rule 1, $p(20) = 0.53$, whereas in Rule 2, $p(20) = 0.42$. The difference is even more apparent in the cumulative rating performance (see Figure 5.2D) where $P(20)$ is almost twice larger in Rule 1 compared to Rule 2.

This illustrates a kind of paradoxical situation in which individuals visit cells with higher values in the competitive Rule 2, but put stars on cells of lower values than in the non-competitive Rule 1. However, as we delve deeper into Section 5.4.2, we will see that this apparent contradiction can be rationalized by the different rating strategies that individuals employ. In particular, under the competitive Rule 2, some individuals give low ratings to cells with a high value, presumably to withhold information by avoiding reporting these cells to the other members of their group.

5.3. Collective Behavior

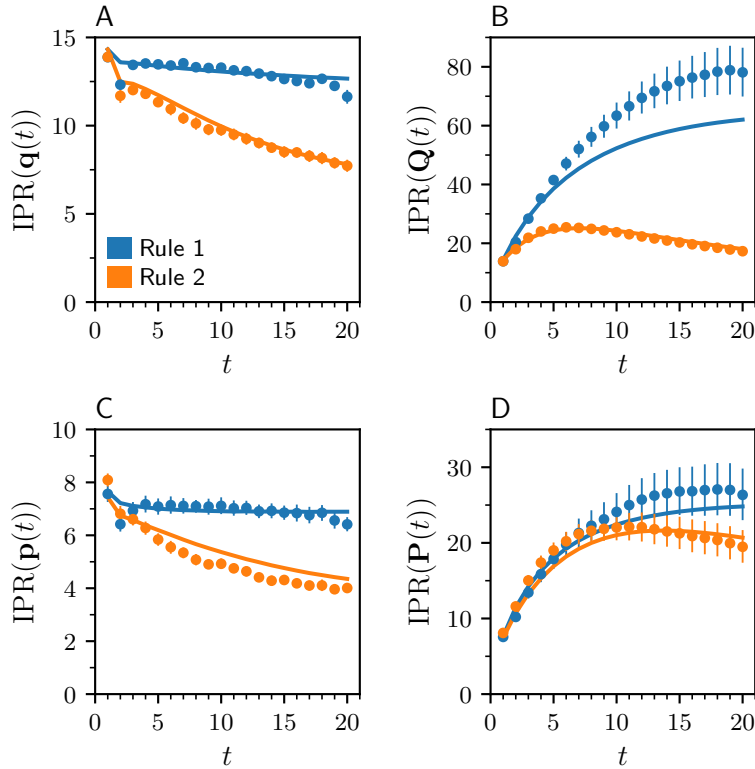


Figure 5.4: **Inverse participation ratio.** Inverse participation ratio of (A) the instantaneous fraction of visits $IPR(\mathbf{q}(t))$, (B) the cumulative fraction of visits $IPR(\mathbf{Q}(t))$, (C) the instantaneous fraction of stars $IPR(\mathbf{p}(t))$, and (D) the cumulative fraction of stars $IPR(\mathbf{P}(t))$ as a function of the round t . The non-competitive Rule 1 is in blue and the competitive Rule 2 is in orange. The dots are the experimental data, and the solid lines are the predictions of the model.

Inverse participation ratio

The Inverse Participation Ratio (IPR) of a distribution is a measure of its spread (see Definition 3.7). Therefore, the quantity $IPR(\mathbf{q}(t))$ (Figure 5.4A) denotes the effective number of cells visited by every player during round t . In the first round, both IPR values are equal because individuals make random choices. However, as the game progresses, individuals visit significantly more cells in each round in Rule 1 than in Rule 2. For example, in the final round, out of a maximum of 15 different cells to explore, they visit about 12 in Rule 1, while they visit only 8 in Rule 2.

Now, looking at $IPR(\mathbf{Q}(t))$ (Figure 5.4B), we observe that the effective number of cells visited up to round t is significantly larger in Rule 1. At the end of the experiment, this number is even 4.5 times larger in Rule 1 with 78 compared to 17 in Rule 2. Notably, in Rule 2, the effective number of cells visited begins to decrease after round 6. Altogether, these results suggest that in Rule 2, participants not only explore a more limited set of cells within the same round, but also tend to revisit the same cells over multiple rounds.

Shifting our focus to the inverse participation ratio of the fraction of ratings (Fig-




Figure 5.5: **Fidelity.** Fidelity of the hidden values of the grid \mathbf{V} to (A) the visits $F(\mathbf{V}, \mathbf{Q}(t))$, and (B) the ratings $F(\mathbf{V}, \mathbf{P}(t))$ as a function of the round t . The non-competitive Rule 1 is in blue and the competitive Rule 2 is in orange. The dots are the experimental data, and the solid lines are the predictions of the model.

ures 5.4C and 5.4D), we observe that individuals allocate stars to more cells in each round in Rule 1 compared to Rule 2. This is obviously a consequence of the fact that individuals visit more cells in Rule 1 than in Rule 2, although we observe a smaller difference between the two rules in Figure 5.4D (ratings) than in Figure 5.4B (visits), presumably because some participants also give high ratings to low-value cells in Rule 2, as confirmed hereafter.

Fidelity

The fidelity (see Definition 3.8) of visits and ratings (Figure 5.5) increases with the round, suggesting that the correlation between participants' visits or ratings and cell values increases as they explore more and more cells. We also observe that the fidelity is higher in Rule 1, which is partly due to the fact that the participants explore and rate more cells in this non-competitive condition, while we will see that they revisit more often high-value cells in Rule 2. It is interesting to note that, in the non-competitive Rule 1, the fidelity of the visits and the ratings are similar. However, in the competitive Rule 2 the visits are much more faithful to the value than the ratings. This is due to the presence of individuals who visit good cells but rate them poorly. Finally, since the scores are still increasing at the end of the game, one can imagine that the longer the game, the better the correlation (see Figure 5.18).

5.4 Individual Behavior

In this section, I will focus on the characterization of the behavior of each individual, as the previous section already led us to anticipate different behaviors in the two rules. In a game, each individual performs two different types of actions: visiting a cell and rating that cell. In a first part, I will focus on how they select a cell to visit. In a second part, I will show how they choose the rating to give to the visited cells based on their value. Finally, in a third part, I will classify individuals into behavioral profiles based on their

5.4. Individual Behavior

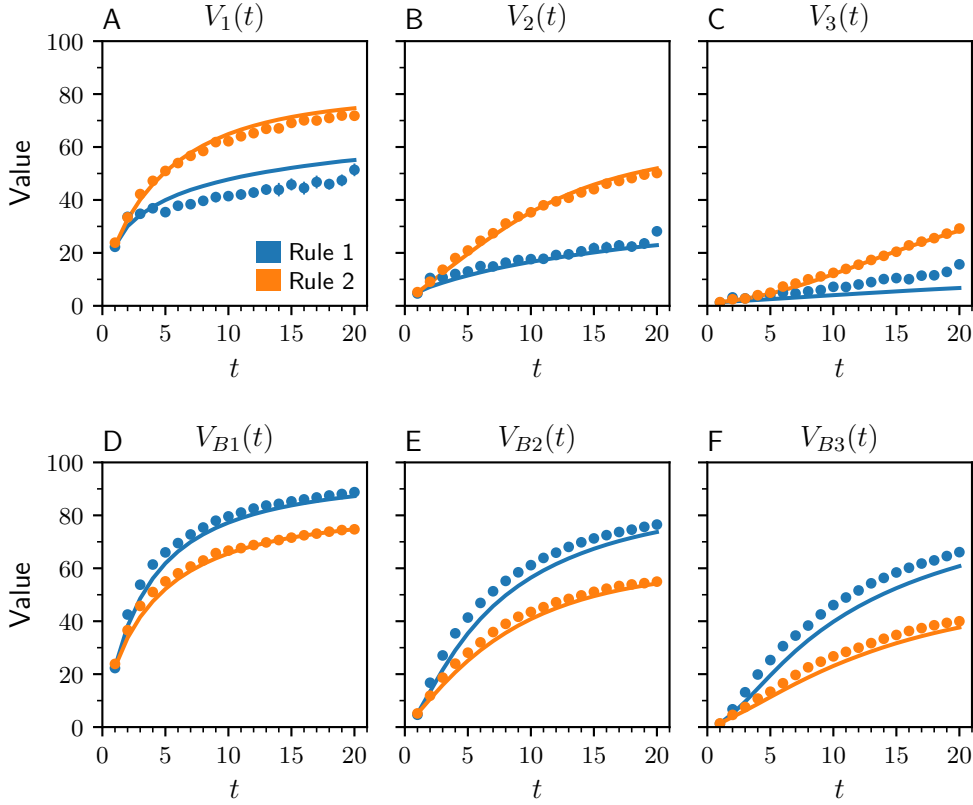


Figure 5.6: **Values of visited cells.** (A–C) $V_1(t)$, $V_2(t)$, $V_3(t)$ are respectively the value of the first-best cell, second-best cell, and third-best cell visited by the participants at round t . (D–F) $V_{B1}(t)$, $V_{B2}(t)$, $V_{B3}(t)$ are respectively the value of the first-best cell, second-best cell, and third-best cell visited by the participants up to round t . The non-competitive Rule 1 is in blue and the competitive Rule 2 is in orange. The dots are the experimental data, and the solid lines are the predictions of the model.

rating behavior.

5.4.1 Visiting cells

Values of visited cells

Let us examine the values associated with the cells opened by the participants in each round (see Definition 3.11 and Figures 5.6A to 5.6C). In both conditions, the average values of the three cells increase with round t , indicating that players are increasingly uncovering cells with higher values. Moreover, on average, individuals in Rule 2 tend to open cells with higher values compared to those in Rule 1, which reflects the competitive nature of Rule 2.

To further illustrate the contrast in exploration strategies between the two rules, we can compare the mean cell values opened at each round to the best values uncovered since the beginning of the game (see Definition 3.12 and Figures 5.6D to 5.6F). In Rule 2, we observe that $V_i(t) \simeq V_{Bi}(t)$ for $i = 1, 2, 3$, meaning that participants consistently revisit the




Figure 5.7: **Probability of finding the cells with the highest values.** (A) Probability to find the best cell, of value 99. (B) Mean probability to find one of the four cells whose values are 86 ($\times 2$), 85, or 84. (C) Mean probability to find one of the four cells whose values are 72 ($\times 2$) or 71 ($\times 2$). The non-competitive Rule 1 is in blue and the competitive Rule 2 is in orange. The dots are the experimental data, and the solid lines are the predictions of the model. The black dashed and dotted lines correspond to the expected probabilities of two different visit strategies: cells chosen randomly (see Equation (3.19)), and cells chosen sequentially (see Equation (3.20)), respectively.

best cells they found. In contrast, in Rule 1, we find that $V_i(t) \simeq V_{B_i}(t)/2$ for $i = 1, 2, 3$, highlighting a distinct pattern where individuals are less inclined to revisit the high-value cells they previously found. Figures 5.6D to 5.6F also show that, in Rule 1, individuals found cells with significantly higher values compared to Rule 2.

These results explain the apparent paradox observed in the previous section. In the competitive Rule 2, players tend to replay the very best cells discovered so far, resulting in a better score than in Rule 1. However, the lesser exploration prevents them from discovering cells as high as those in Rule 1. In Rule 1, the absence of incentives to revisit high-value cells results in a more exploratory behavior, which is corroborated by both the higher value of the inverse participation ratio and the higher fidelity of ratings to values.

Finding the highest-value cells

To get a clearer insight into the phenomenon described in the previous section, Figure 5.7 shows the probabilities associated with the discovery of the highest-value cells (see Definition 3.14). The probabilities increase with each successive round, with Rule 2 showing a comparatively slower increase rate, as expected. In particular, at the end of the game, the probability of uncovering the cell with the highest value (99) is about four times higher in Rule 1 ($\simeq 0.40$) than in Rule 2 ($\simeq 0.10$).

Figure 5.7 also show the expected results of two simple visit strategies introduced in Definition 3.14: random visits and sequential visits. The latter represents the highest achievable probabilities in the absence of ratings. A curve above this threshold would indicate information sharing among participants, where those who initially found high-value cells share this information, and subsequently other participants visit these cells. This

5.4. Individual Behavior

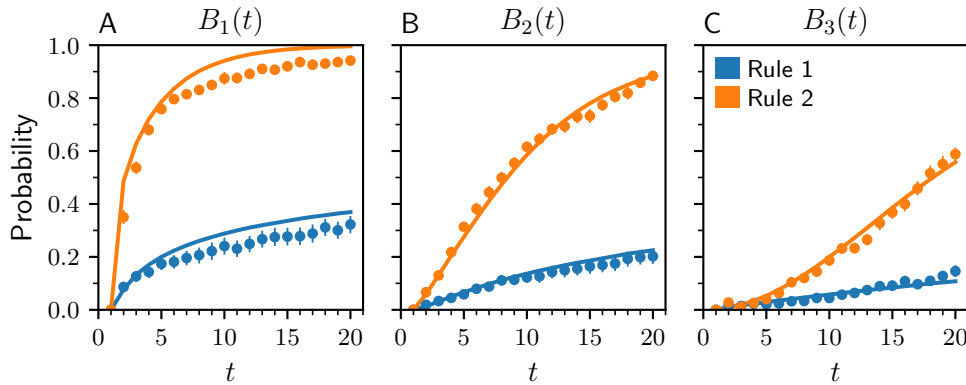


Figure 5.8: **Probability of revisiting cells visited in the previous round.** (A–C) Probability $B_1(t)$, $B_2(t)$, $B_3(t)$ to revisit the first-best cell, the second-best cell, and the third-best cell visited in the previous round as a function of the round t for $t > 1$. The non-competitive Rule 1 is in blue and the competitive Rule 2 is in orange. The dots are the experimental data, and the solid lines are the predictions of the model.

cooperative scenario is observed in Rule 1. In Rule 2, the curve falls below the probability predicted by the two simple strategies. This could indicate a lack of communication among participants, which will be explored in the following section. Alternatively, it could indicate a strategy in which participants over-visit cells that are deemed sufficiently high in value to increase their scores. This strategy may cause them to avoid exploring new cells and consequently fail to find cells with higher values.

Revisiting the high-value cells

Our analysis so far (and in particular, Figure 5.6) suggests that, in Rule 2, individuals frequently revisit their best cells. To quantify this behavior, Figure 5.8 shows the probability of revisiting cells already visited in the previous round (see Definition 3.13). Note that the probability values are particularly high in Rule 2. In the final round, participants revisit their first-best, second-best, and third-best cells from the previous round with probabilities of 0.94, 0.88, and 0.59, respectively. Conversely, in Rule 1 individuals have comparatively lower probabilities. In the final round, individuals revisit their first-best, second-best, and third-best cells from the previous round with probabilities of 0.32, 0.20, and 0.15, respectively. This confirms the observation that in Rule 1 participants predominantly explore new cells, whether they are white or colored. In contrast, participants in Rule 2 tend to explore fewer cells and revisit cells they have already opened.

5.4.2 Rating cells

To better understand how the cells are rated, we have measured the mean rating $R(V)$ (see Definition 3.15), which is the average number of stars used to rate a cell as a function of its value. Figure 5.9 shows a significant contrast in the way individuals rate cells between each condition. In Rule 1, the mean rating increases almost linearly with the value of the cell. On average, individuals give 1 star to the cells with low values and 4.3 stars to the




Figure 5.9: **Mean rating of the group.** Average number of stars used to rate cells by every player in a game as a function of the value V of the cell. The non-competitive Rule 1 is in blue and the competitive Rule 2 is in orange. The dots are the experimental data, and the solid lines are the predictions of the model.

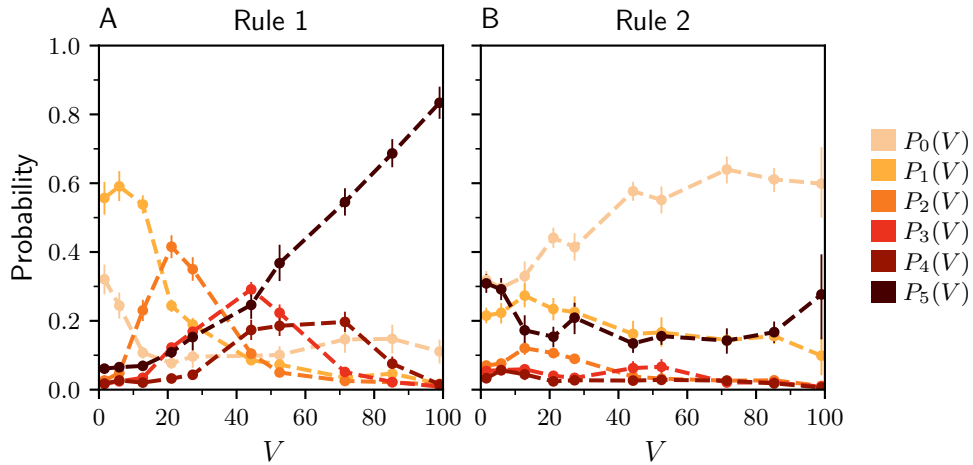


Figure 5.10: **Rating probability.** Probability of rating a cell as a function of its value V with $s = 0, 1, \dots, 5$ stars $P_s(V)$ in (A) the non-competitive Rule 1 and (B) the competitive Rule 2.

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 mean rating decreases to reach a plateau at about 1.5 stars for values higher than $V = 25$. Thus, a cell with value between 35 and 99 receives similar mean ratings regardless of its value. This phenomenon suggests that in Rule 2, many participants adopt a non-cooperative/deceptive rating strategy, which effectively makes the information conveyed by the digital trace less discriminating. Overall, these results indicate that individuals give a much fairer rating to the cells they visit in Rule 1, as the examination of the fidelity has previously revealed (see Figure 5.5B).

Let us now examine the ratings in more detail by analyzing the distribution of ratings for each value. Figure 5.10 shows the probability of assigning s stars to a cell with value V (see Definition 3.17). One can observe two distinct patterns. In Rule 1, individuals use the full range of ratings, with a predominant tendency to assign a majority of 5-star ratings

5.4. Individual Behavior

to the cells of high values. Conversely, in Rule 2 one can observe a very different trend where all probabilities remain nearly constant for all ratings, except for the probability of assigning 0 stars, which increases with the value of the cell. In addition, the probabilities of assigning 2, 3, or 4 stars are close to zero, suggesting that in Rule 2 individuals predominantly use the 0, 1, and 5-star ratings, while neglecting the intermediate rating options.

5.4.3 Behavioral profiles of individuals

To gain a better understanding of individual rating behaviors, we have to study how each individual rates cells. This can be achieved by examining the mean rating of each individual in each game. For representation reasons, individual mean ratings for each game will not be displayed due to the large number of games. However, Figures 5.11 and 5.12 display the mean rating of each participant over the ten games played in groups of five. This gives less variety as individuals can change their rating strategy in each game, but it provides a good overview of the situation. One can observe that in both rules there exists a wide variety of behaviors. Some individuals rate cells somewhat proportionally to their value, some rate cells almost independently of their value, and some others give ratings somewhat oppositely proportional to the cell values.

To better quantify the behaviors of each individual, we have fitted the mean rating $R(V)$ of each individual using a linear function (see Section 3.4.2 for more details)

$$R(V) = u_0 + 5 \frac{V}{V_{\max_1}} u_1, \quad (5.1)$$

where u_0 is the intercept, u_1 is the slope of the line, and $V_{\max_1} = 90$ is the maximum cell value of the grid.

The combination of every pair of parameters u_0 and u_1 for all participants is shown in Figure 5.13, along with the respective distributions of u_0 and u_1 . The inspection of the distribution of u_1 suggests the existence of 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 .

This classification leads to three behavioral profiles:

- Individuals with $u_1 \geq u_{\text{neu-col}}$ rate cells with a rating that increases with the cell values, i.e., they rate cells whose values are the lowest (resp. whose values are the highest) with a small number of stars (resp. a large number of stars; see Figure 5.14A). Hereafter, we will dub these individuals as *collaborators*, since their rating strategy helps the other members of their group to identify the best cells. This altruistic behavior can come at a cost in Rule 2, since helping the other members of their group may negatively impact the final ranking of collaborators, and hence their monetary reward.
- Individuals with $u_{\text{def-neu}} \leq u_1 < u_{\text{neu-col}}$ rate cells with almost the same number of stars regardless of their values (see Figure 5.14B). Since the ratings of these individuals do not provide any distinctive information to the other group members,




Figure 5.11: **Individual mean rating on all games for Rule 1.** Average number of stars used to rate cells by one player in every game of an experimental session as a function of the cell's value V . Each of the rectangles corresponds to the behavior of a single individual aggregated on the 10 games. The x -axis is the cell's value and goes from 0 to 100 and the y -axis is the number of stars used by the individual to rate a cell of a given value and goes from 0 to 5. The dots are the experimental data, and the line is a linear fit of these data with the function given by Equation (5.1). Individuals are sorted from left to right and from top to bottom according to the value of the slope u_1 . The color corresponds to the behavioral profile: green for collaborators, brown for neutrals, and red for defectors.




Figure 5.12: **Individual mean rating on all games for Rule 2.** Average number of stars used to rate cells by one player in every game of an experimental session as a function of the cell's value V . Each of the rectangles corresponds to the behavior of a single individual aggregated on the 10 games. The x -axis is the cell's value and goes from 0 to 100 and the y -axis is the number of stars used by the individual to rate a cell of a given value and goes from 0 to 5. The dots are the experimental data, and the line is a linear fit of these data with the function given by Equation (5.1). Individuals are sorted from left to right and from top to bottom according to the value of the slope u_1 . The color corresponds to the behavioral profile: green for collaborators, brown for neutrals, and red for defectors.

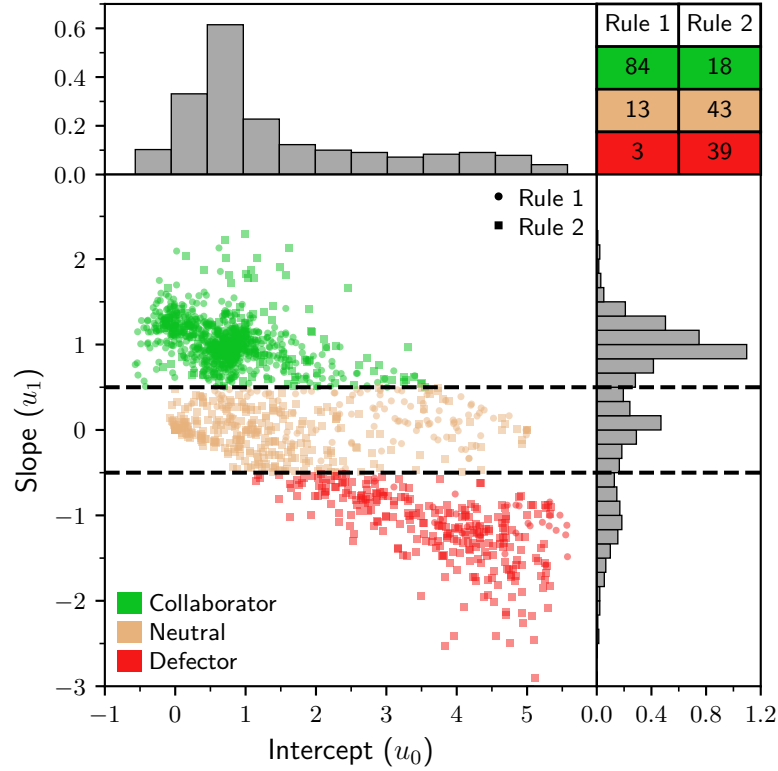


Figure 5.13: **Behavioral profiles of individuals.** (Bottom-left) Scatter plot of the values of the two parameters u_0 and u_1 of the linear function (Equation (5.1)) used to fit each participant's mean rating as a function of the value of the visited cells. In the non-competitive Rule 1, individuals 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) Distribution of the values of u_0 . (Bottom-right) Distribution of the values of u_1 . (Top-right) The table gives the percentage of individuals for each of the behavioral profiles.

5.4. Individual Behavior

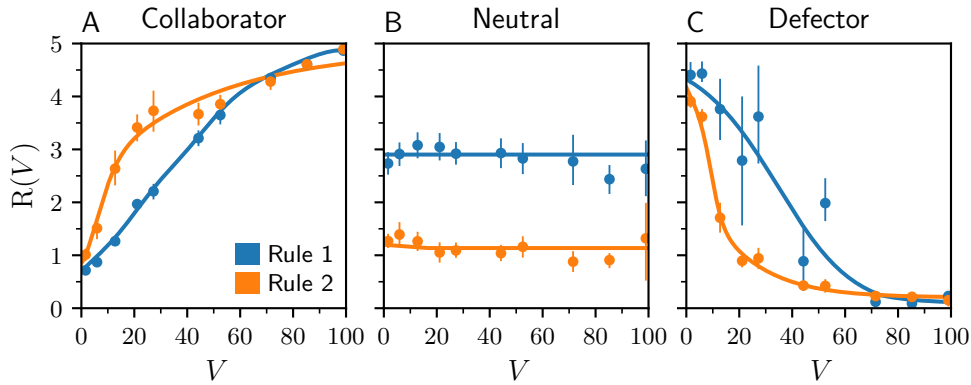


Figure 5.14: **Mean rating of each behavioral profile.** Average number of stars used to rate cells as a function of the cell’s value V for (A) collaborators, (B) neutrals, and (C) defectors. The non-competitive Rule 1 is in blue and the competitive Rule 2 is in orange. The dots are the experimental data, and the solid lines are the predictions of the model.

we will dub them as *neutrals*. Note that these neutral individuals do not form a homogeneous group. Indeed, some of them with u_0 close to 0 always give 0 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 5 always give a large number of stars or even 5 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 Figure 5.13. 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.

- 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 (resp. a large number of stars) to the cells whose values are the highest (resp. whose values are the lowest; see Figure 5.14C). We will call these individuals *defectors*, 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. In addition, they also decide not to share the position of the best cells they have found, by giving them low ratings, and hence not marking them on the grid.

Note that the two thresholds separating these three behavioral classes are close to the corresponding thresholds obtained using Ward’s clustering method on the slope parameter, u_1 . We have also developed a machine learning approach where a network is trained with games generated by our model, again leading to similar results. In the rest of the thesis, we will fix these thresholds to ± 0.5 , which will also allow us to compare the fraction of the three behavioral profiles in different experimental conditions (see Chapters 6 and 7). Also note that the couple of value $u_0 = 0$ and $u_1 = 1$ (respectively, $u_0 = 5$ and $u_1 = -1$) corresponds to an ideal collaborator (respectively, an ideal defector) rating cells strictly proportionally (respectively, oppositely proportionally) to their value and using the full

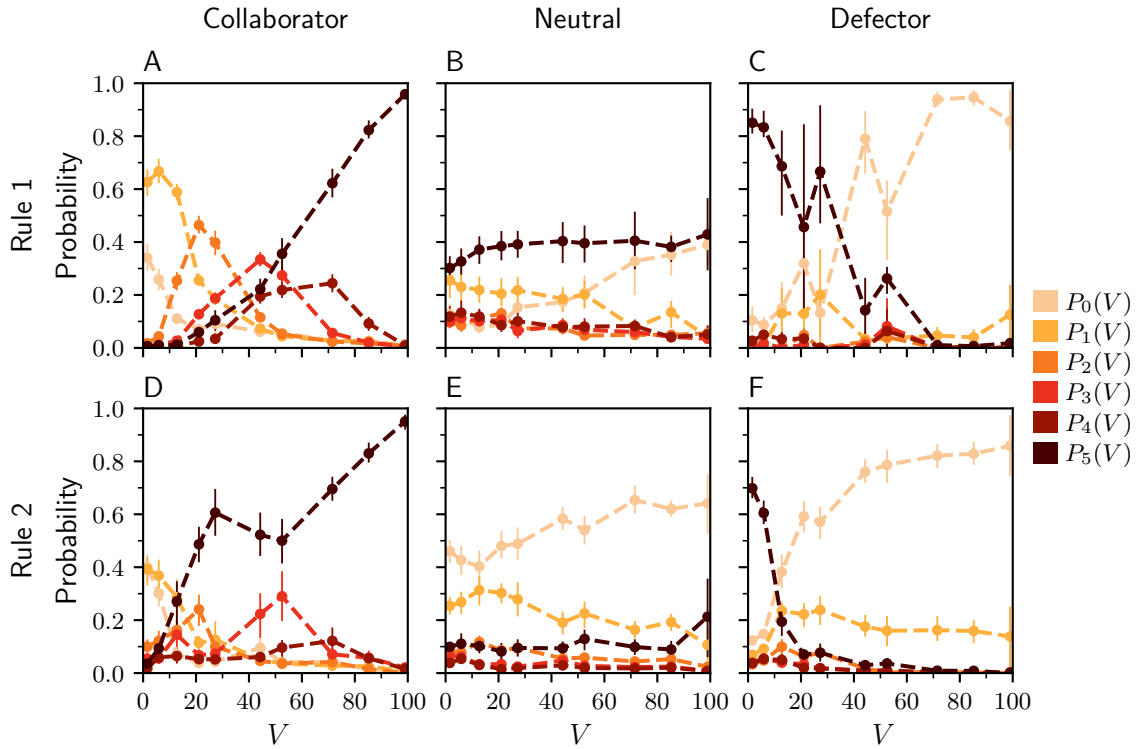


Figure 5.15: **Rating probability of each behavioral profile.** Probability $P_s(V)$ of rating a cell of value V with $s = 0, 1, \dots, 5$ stars for each behavioral profile in Rule 1 (top row) and Rule 2 (bottom row).

rating scale between 0 and 5 stars (see Section 3.4.2). Hence, it certainly makes sense to adopt fixed thresholds for u_1 , equal to half of the slopes for ideal collaborators and defectors.

Looking at the percentages of each behavioral profile in both rules (see Table 5.1), one observes that in Rule 1, most of the individuals are collaborators (84%), with few neutrals (13%), and almost no defectors (3%). This results in a high-quality shared information in the game, with the presence of a high fraction of collaborators ensuring that the cell colors on the grid accurately reflect their value (also leading to a high fidelity). However, in Rule 2, neutrals (43%) and defectors (39%) dominate, and we only observe a small fraction of collaborators (18%), resulting in a low-quality shared information.

	Rule 1	Rule 2
Collaborator	84 ± 3	18 ± 2
Neutral	13 ± 2	43 ± 3
Defector	3 ± 2	39 ± 4

Table 5.1: **Percentages of each behavioral profile.** Percentages of each behavioral profile: collaborator, neutral, and defector for both rules.

Figure 5.15 shows that collaborators mostly rate cells whose values are less than 20 with

5.5. Model Predictions




Figure 5.16: **Distribution of behavioral profiles as a function of the rank.** Fraction of individuals with each behavioral profile (collaborator, neutral, and defector) found at ranks $r = 1, 2, \dots, 5$ in Rule 2. The colored bars correspond to experimental data for each behavioral profile: collaborator (green), neutral (brown), and defector (red) and the black lines are the predictions of the model. The horizontal dashed lines are the proportion of individuals of each behavioral profile in all experiments (null model).

1 star, while the cells whose values are greater than 80 are rated with 5 stars. By contrast, for the neutral individuals, the probability of rating a cell with a given number of stars does not depend on the cell value. Finally, 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 and neutrals have access to more information than other group members. In fact, they benefit from collaborators who give high ratings to cells with high values. At the same time, defectors strategically withhold their knowledge about the best cells they have found by refraining from marking such cells, and neutrals do not provide discriminating information to other members. Thanks to this asymmetric information [Balakrishnan and Koza, 1993], defecting or, to a lesser extent, neutral behavior can be advantageous in a competitive environment. Indeed, defectors and neutrals have a higher probability of having the highest score in their group (see Figure 5.16). However, in the absence of competition, there is no benefit to deception and one should expect fewer defectors. This is what we observe in our experiments, where Table 5.1 shows that almost every participant adopts cooperative behavior in Rule 1, while there is a large fraction of defectors in Rule 2.

5.5 Model Predictions

This section applies the model introduced in Chapter 4 to the experiment in Rule 1 and Rule 2. The model enhances our understanding of human behavior in the experiments and enables predictions of human behavior in unstudied situations. Additionally, the model will be used to design agents that optimize specific quantities, providing valuable insights into the game's underlying mechanisms. In Chapter 6, we will even use the model to

control four bots playing against a human participant.

5.5.1 Modeling human behavior

Let us first model the participants' behavior by introducing "Mimic agents" that reproduce human behaviors within the game as observed in the two experimental conditions. To create these agents, one has to define two key components: the visit strategy and the rating strategy.

This study identified three distinct behavioral profiles (collaborator, neutral and defector), each requiring a unique rating strategy. The probability distribution of ratings, as shown in Figure 5.10, indicates that not all ratings are equally prevalent. Except for collaborators in Rule 1, participants mainly assign ratings of either 0 or 5 stars. Ratings of 1, 2, 3, or 4 stars are less common and have similar probabilities. To keep the model as simple as possible in except for the collaborators in Rule 1, the probabilities of rating a cell with 1 to 4 stars are set equal. This is achieved through a probabilistic normalization condition, which requires that the sum of probabilities for each value of V from 0 to 5 equals 1. Specifically, for $s = 1, 2, 3, 4$:

$$P_s(V) = P_{1234}(V) = \frac{1}{4}(1 - P_0(V) - P_5(V)). \quad (5.2)$$

The probabilities of assigning ratings of 0 and 5 stars are determined by fitting experimental values using a sigmoid function (see Equation (4.4)) for collaborators and defectors, or a linear function (see Equation (4.6)) for neutrals, while maintaining the observed mean rating.

The $P_{1234}(V)$ approximation does not apply to collaborators in Rule 1, since they use the entire rating scale to rate cells proportionally to their values. In this case, we fitted the experimental values using a normal function (see Equation (4.5)) for each rating, except for the zeroth one, which we set as $P_0(V) = 1 - \sum_{s=1}^5 P_s(V)$.

The obtained parameter values for both rules are reported in Table 5.2, and the resulting fits are presented in Figure 5.17.

After establishing a rating strategy, the next step is to determine the parameters for the visit strategy. We conducted simulations with five identical agents, each of them using a rating strategy randomly drawn from collaborative, neutral, and defector profiles based on the proportions observed in the experiments (see Table 5.1). The Monte Carlo method outlined in Chapter 4 was used to iteratively adjust the parameters and minimize the error (see Equation (4.7)) between experimental and simulated observables.

The error metrics considered encompass a range of observables: $q(t)$, $Q(t)$, $p(t)$, $P(t)$, $\text{IPR}(\mathbf{q}(t))$, $\text{IPR}(\mathbf{Q}(t))$, $\text{IPR}(\mathbf{p}(t))$, $\text{IPR}(\mathbf{P}(t))$, $F(\mathbf{Q}(t), \mathbf{V})$, $F(\mathbf{P}(t), \mathbf{V})$, $V_1(t)$, $V_2(t)$, $V_3(t)$, $B_1(t)$, $B_2(t)$, and $B_3(t)$. The error is equal to $\Delta = 0.014$ in Rule 1 and to $\Delta = 0.005$ in Rule 2. The resulting optimized parameters are detailed in Table 5.3.

5.5.2 Comparison between theoretical and experimental results

We consider groups of five Mimic agents, reproducing the behaviors of human collaborators, neutrals, and defectors. Their behavioral profiles are drawn according to the

5.5. Model Predictions

	d'_s	e'_s	f'_s
$s = 1$	0.65	6.6	5.83
$s = 2$	0.46	25.9	6.30
$s = 3$	0.36	43.8	4.79
$s = 4$	0.30	61.1	4.07
$s = 5$	0.96	102.4	2.01

(A) Collaborator (Rule 1)

	c_s	d_s	e_s	f_s
$s = 0$	1113.4	1113.3	-84.5	-4.75
$s = 5$	-1051.9	1052.8	-304.5	1.24

(B) Collaborator (Rule 2)

	c''_s	f''_s
$s = 0$	0.09	0.30
$s = 5$	0.25	0.30

(C) Neutral (Rule 1)

	c''_s	f''_s
$s = 0$	0.45	0.17
$s = 5$	0.09	0.17

(D) Neutral (Rule 2)

	c_s	d_s	e_s	f_s
$s = 0$	0.50	0.45	39.4	3.86
$s = 5$	0.46	0.52	26.9	-3.11

(E) Defector (Rule 1)

	c_s	d_s	e_s	f_s
$s = 0$	0.45	0.46	14.8	7.34
$s = 5$	0.39	0.38	9.8	-18.49

(F) Defector (Rule 2)

Table 5.2: **Rating strategy parameters of Mimic agents.** Parameters values used for the rating strategy (see Equations (4.4) to (4.6)) for the Mimic agents (collaborator, neutral, and defector) in both rules.

	$P^E(c, t)$		$B_1(t)$		$B_2(t)$		$B_3(t)$	
	ε	α	a_1	b_1	a_2	b_2	a_3	b_3
Rule 1	0.78	0.89	57.6	2.19	25.0	2.29	1.4	2.64
Rule 2	0.69	1.32	-8.4	1.55	-4.1	2.11	-0.2	2.33

Table 5.3: **Visit strategy parameters of Mimic agents.** Parameters values used for the visiting strategy (see Equations (4.1) and (4.2)) for the Mimic agents.

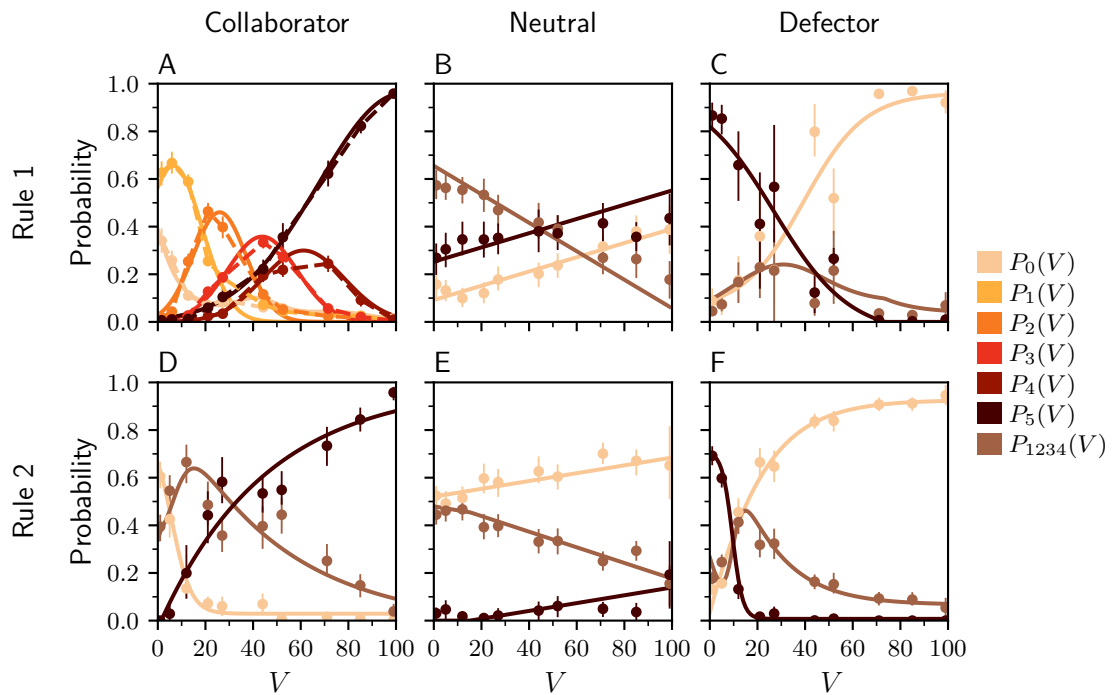


Figure 5.17: **Rating strategy of Mimic agents.** Probability of rating a cell with s stars ($P_s(V)$), for the collaborators, neutrals, and defectors, and for the two rules. Except for collaborators in Rule 1, the probabilities of rating a cell of value V with 1 to 4 stars have been averaged in $P_{1234}(V)$. The dots are the experimental data, and the solid lines represent the rating strategy in the model for collaborator, neutral, and defector agents.

5.5. Model Predictions

corresponding fraction observed in the experiment (see Table 5.1). All graphs were based on 1,000,000 simulations, resulting in negligible error bars on the scale of the presented graphs.

Figures 5.1, 5.2, 5.4 and 5.5 show 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 (Figure 5.6), along with the probability to revisit each of these three best cells at the next turn (Figure 5.8). The model reproduces fairly the fraction of collaborators, neutrals, or defectors according to their rank at the end of the experiment (Figure 5.16).

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.

5.5.3 Impact of the number of rounds and group size on individual performance and collective dynamics

Figure 5.18 shows that after 100 rounds, instead of 20 rounds, the normalized score of individuals and groups has increased by 60% in Rule 2. Beyond round 50, the values of the observables used to quantify the dynamics of collective exploration and ratings begin to saturate. From one round to another, the Mimic agents revisit almost exclusively the same cells whose values are very high. At the end of the 100 rounds, in Rule 2 the value of their best cell is $V_1(100) \simeq 84$, and the agents revisit their best cell with a probability $B_1(100) \simeq 1$.

Figure 5.19 shows the impact of group size on the scores of individuals and groups, and the dynamics of collective exploration and ratings. We compare the simulation results obtained with groups of 5 Mimic agents exploring a grid of 225 (15×15) cells and groups of 20 Mimic agents exploring a grid four times larger, 900 cells (30×30). These larger grids were obtained from the combination of four identical grids of 225 cells so that the proportion of each cell value does not change. For instance, in a grid of 900 cells, there are four cells with a value of 99, but their proportion ($1/225$) is the same as in the smaller grids. The dynamics of the inverse participation ratio (IPR) of $\mathbf{p}(t)$, $\mathbf{P}(t)$, $\mathbf{Q}(t)$, and $\mathbf{Q}(t)$ reveal that large groups do not visit four times more cells than small groups, but instead, they concentrate their visits on a few cells with high values. Individuals also have a higher probability of finding the cells with the best values. However, despite these differences, the score remains unchanged. Finally, in Rule 1, the probability that individuals find the best cells at the end of an experiment is much larger in groups of 20 Mimic agents. Altogether, these results suggest that cooperation induced by stigmergic interactions and the way individuals use the traces resulting from past actions increase with group size.

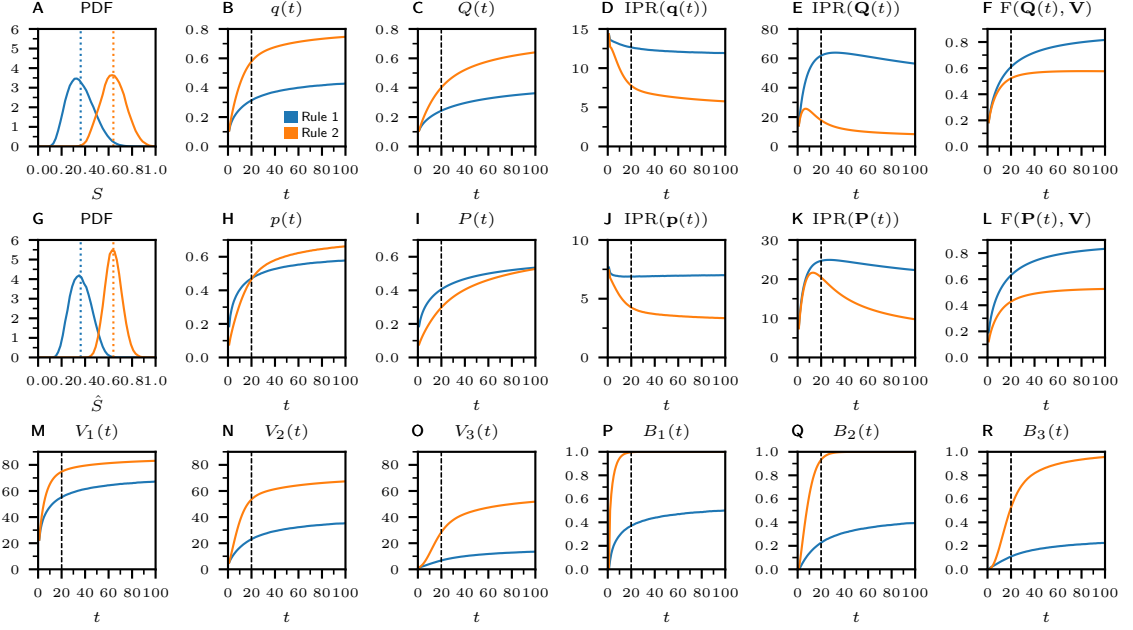


Figure 5.18: Impact of the number of rounds on individual performance and collective dynamics. Simulations with five Mimic agents over 100 rounds. The non-competitive Rule 1 is in blue and the competitive Rule 2 is in orange. The vertical dotted line at $t = 20$ corresponds to the final round used in the experiments with humans. (A) Probability distribution function (PDF) of the normalized individual scores S , and (G) normalized group scores \hat{S} . The dotted vertical lines are the mean normalized scores. (B) Average value of the cells visited at round t , $q(t)$ and (C) up to round t , $Q(t)$. (H) Average value of the cells visited weighted by their ratings at round t , $p(t)$ and (I) up to round t , $P(t)$. (D) and (E) Inverse participation ratio of the visits, $\text{IPR}(\mathbf{q}(t))$ and $\text{IPR}(\mathbf{Q}(t))$. (J) and (K) Inverse participation ratio of the ratings, $\text{IPR}(\mathbf{p}(t))$ and $\text{IPR}(\mathbf{P}(t))$. (F) Fidelity to the cell value distribution of the distribution of visits, $F(\mathbf{Q}(t), \mathbf{V})$, and, (L) of ratings, $F(\mathbf{P}(t), \mathbf{V})$. (M–O) $V_1(t)$, $V_2(t)$, $V_3(t)$ 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 . (P–R) 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 $t > 1$.

5.5. Model Predictions

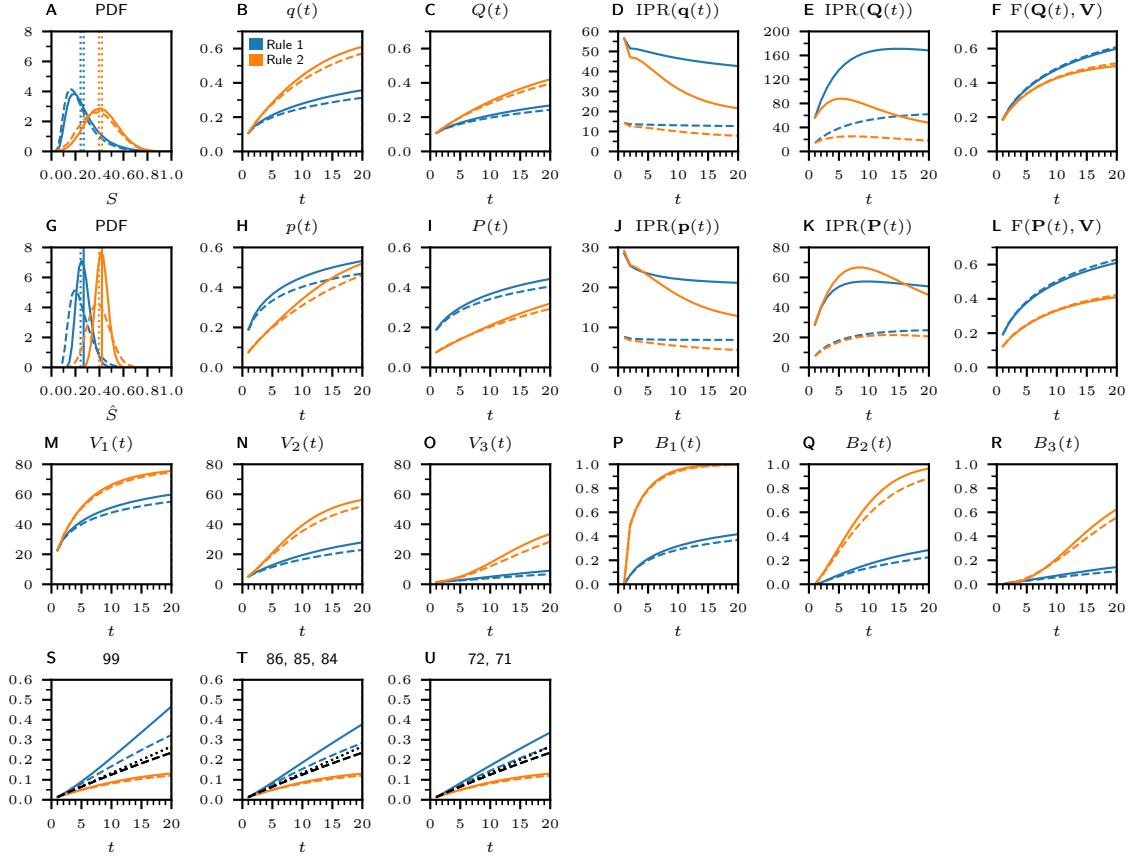


Figure 5.19: **Impact of the group size on the collective performance and the dynamics of collective exploration and ratings.** Dashed lines correspond to simulations with five Mimic agents exploring a grid with 225 (15×15) cells, as used in the experiments with humans. Solid lines correspond to simulations with 20 Mimic agents exploring a grid 4 times larger, with 900 (30×30) cells. The non-competitive Rule 1 is in blue and the competitive Rule 2 is in orange. (A) Probability distribution function (PDF) of the normalized individual scores S , and (G) normalized group scores \hat{S} . The dotted vertical lines are the mean normalized scores. (B) Average value of the cells visited at round t , $q(t)$ and (C) up to round t , $Q(t)$. (H) Average value of the cells visited weighted by their ratings at round t , $p(t)$ and (I) up to round t , $P(t)$. (D) and (E) Inverse participation ratio of the visits, $\text{IPR}(\mathbf{q}(t))$ and $\text{IPR}(\mathbf{Q}(t))$. (J) and (K) Inverse participation ratio of the ratings, $\text{IPR}(\mathbf{p}(t))$ and $\text{IPR}(\mathbf{P}(t))$. (F) Fidelity to the cell value distribution of the distribution of visits, $F(\mathbf{Q}(t), \mathbf{V})$, and, (L) of ratings, $F(\mathbf{P}(t), \mathbf{V})$. (M–O) $V_1(t)$, $V_2(t)$, $V_3(t)$ 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 . (P–R) 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 $t > 1$. (S) Probability to find the best cell, of value 99. (T) Probability to find one of the four cells whose values are 86 ($\times 2$), 85, or 84. (U) Probability to find one of the four cells whose values are 72 ($\times 2$) or 71 ($\times 2$).

5.5.4 Impact of the rating strategy on agents' performance and the fidelity of ratings

To better understand the impact of the rating strategy on individual performance, we studied the collective behaviors of groups of five agents having a linear rating strategy (see Section 4.3.1). These agents rate a cell proportionally to its value, V , with $u_0 + u_1 \times 5V/99$ stars, where u_0 and u_1 are respectively the intercept and the slope of the line (see Equation (5.1)). When $u_1 > 0$, the number of stars used to rate a cell increases with its value V (like for a cooperator), while when $u_1 < 0$, the number of stars used to rate a cell decreases with its value V (like for a defector). As u_0 increases, agents use a larger number of stars to rate a cell of a given value. Moreover, the combinations of parameters $u_0 \leq 0$ and $u_1 \leq 0$ correspond to a situation in which the agents rate all cells with 0 stars, as some actual neutrals do in the experiment. Finally, the visit strategies of these agents are the same as those used by the Mimic agents in each of the two conditions, Rule 1 and Rule 2.

Figure 5.20 presents the result of the respective impact of u_0 and u_1 on (i) the average performance of individuals, (ii) the average value of cells visited by the participants weighted by their ratings, and (iii) the fidelity of ratings with respect to cell values, for each condition Rule 1 and Rule 2. We first observe that when $u_0 = 0$, as soon as the agents start rating the cells with a non-zero number of stars, the resulting trace allows them to cooperate and significantly increase their performance, even for very low positive values of u_1 . The results of the simulations also show that the agents get the best scores for negative values of u_0 , which correspond to situations in which there exists a minimum threshold in the value of a cell that triggers the agents to rate that cell (e.g., when $u_0 = -0.5$ and $u_1 = 0.5$ the threshold is at $V = 20$). Moreover, the higher the value of u_0 , the worse the performance of the agents. This results from the fact that in that condition, the agents use a very high number of stars with little discrimination in the ratings for different values of V . The resulting trace left on cells then provides much less information to the agents, leading to a lower level of cooperation and lower performance. Note however that for high values of u_0 (i.e., when $u_0 > 3$) and for weakly negative values of u_1 (i.e., when $-1 < u_1 < 0$), there still exists weak cooperation between the agents. At first glance, this is rather counterintuitive, since for these parameters, agents are classified as neutrals or mild defectors. However, this phenomenon can be explained by the fact that, while the traces left by the agents in the initial rounds may not allow for the identification of cells with higher values, over time, cells with higher values will be revisited more often, resulting in a greater accumulation of marks compared to cells with lower values. Nevertheless, for values of u_1 that are even more negative, indicating strong defection, the tendency of agents to revisit high-value cells is insufficient to counterbalance the negative impact of assigning high ratings to cells with low values, which ultimately leads to decreased performance.

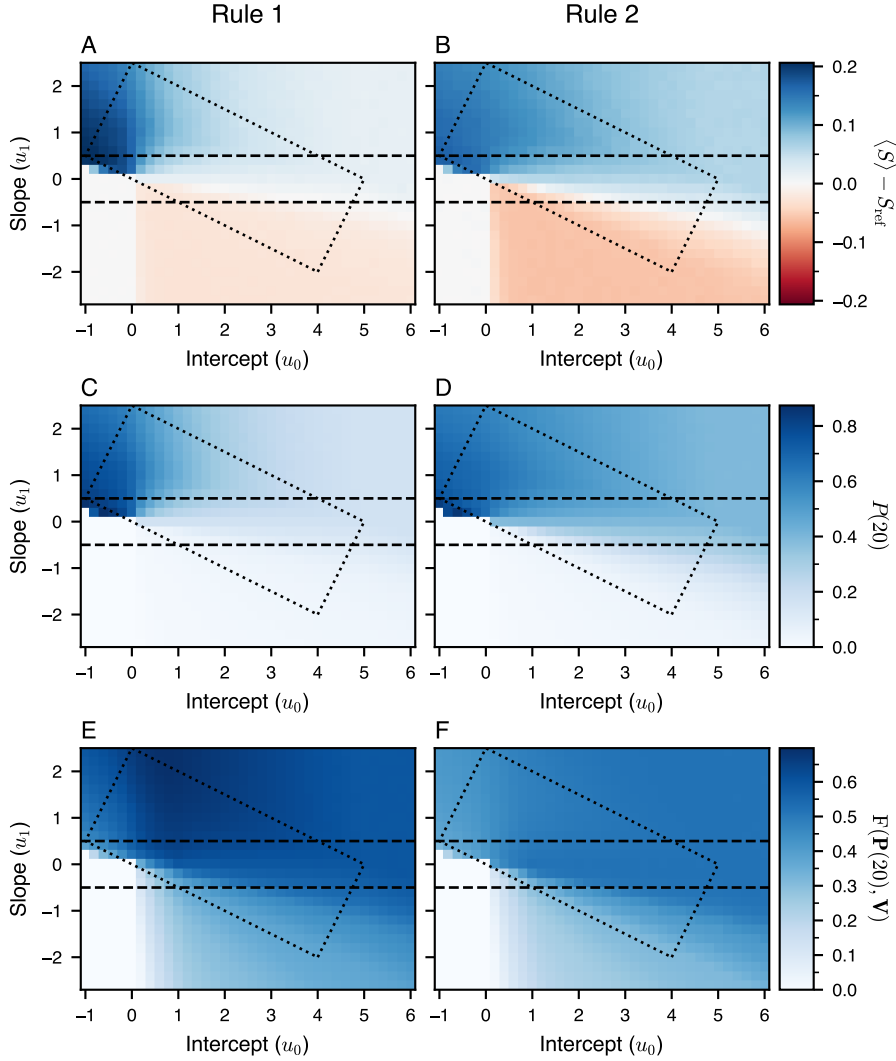


Figure 5.20: **Impact of the rating strategy on agents' performance and the fidelity of rating.** Heatmap for Rule 1 (left column) and Rule 2 (right column) and for different combinations of values of intercept u_0 and slope u_1 of: (A) and (B) the average value of the normalized individual score $\langle S \rangle - S_{\text{ref}}$, (B) and (C) the average value of the cells visited weighted by their ratings at the end of the experiment $P(20)$, and (E) and (F) the average value of the fidelity of ratings with respect to cell values at the end of the experiment $F(\mathbf{P}(20), \mathbf{V})$. Each data point on the heatmap corresponds to the average over 10,000 simulations with five identical agents, defined by their intercept u_0 and slope u_1 . In (A) and (B), S_{ref} is the normalized score obtained with simulations done with $u_0 = 0$ and $u_1 = 0$. Blue (resp. red) corresponds to positive (resp. negative) values, see color bars. The two horizontal lines at $u_{\text{def-neu}} = -0.5$ and $u_{\text{neu-col}} = 0.5$ are the delimitation between the behavioral profiles, and the dotted rectangle represents the rough location of the participants' parameters in the experiment.

	c_s	d_s	e_s	f_s
$s = 0$	0.5	0.95	63.7	-5.17
$s = 5$	0.5	0.78	80.1	5.01

(A) Opt-1

	c'_s	d'_s
$s = 0$	1	0
$s = 5$	0	0

(B) Opt-2

	c_s	d_s	e_s	f_s
$s = 0$	0.45	0.59	16.9	7.34
$s = 5$	0.51	0.55	9.8	-18.48

(C) Opt-3

Table 5.4: **Rating strategy parameters of Optimized agents.** Parameters values used for the rating strategy (see Equations (4.4) to (4.6)) for the Optimized agents: (A) Opt-1, (B) Opt-2, and (C) Opt-3.

	$P^E(c, t)$		$B_1(t)$		$B_2(t)$		$B_3(t)$	
	ε	α	a_1	b_1	a_2	b_2	a_3	b_3
Opt-1	1e-5	1.38	25.0	2.00	18.4	2.03	27.1	2.41
Opt-2	0.58	2.75	-2.4	2.15	4.0	2.54	9.1	2.90
Opt-3	0.82	4.32	22.3	4.86	13.7	3.54	8.3	3.35
Opt-4	1	0	0	0	0	0	0	0

Table 5.5: **Visit strategy parameters of Optimized agents.** Parameters values used for the visit strategy (see Equations (4.1) and (4.2))

5.5.5 Optimization of agents' performance according to specific objectives

We have also exploited our model to find agents that are optimized in different situations. To achieve 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 normalized score S of five identical agents (Opt-1 agents) in the same group and exploiting the same strategy (see Figures 5.21A and 5.22 and Tables 5.4A and 5.5). The inspection of the Opt-1 agents' 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 normalized score ($S \simeq 67\%$) is markedly higher than the normalized score of the human subjects in Rule 2 ($S \simeq 39\%$). Note that, since the five 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 (see Chapter 7, where we present the results of an experiment where two groups directly compete).

We then consider a situation in which we maximize the normalized score of one agent

5.5. Model Predictions

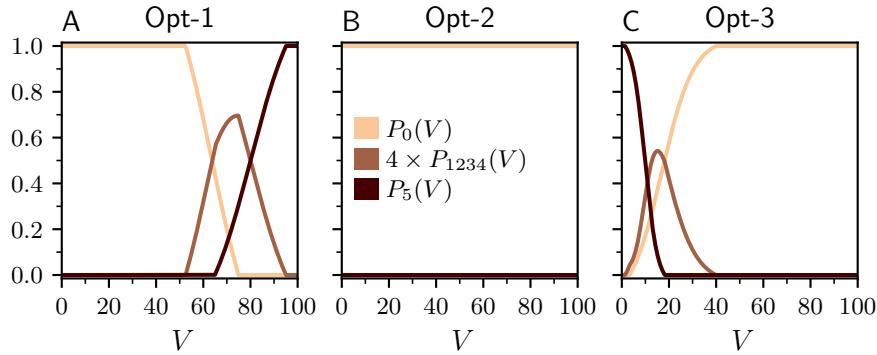


Figure 5.21: **Rating probabilities of Optimized agents.** Probability of rating a cell with 0 stars ($P_0(V)$), 1 to 4 stars ($P_{1234}(V)$) and 5 stars ($P_5(V)$) as a function of its value V , for (A) Opt-1, (B) Opt-2, and (C) Opt-3 agents.

competing with four Mimic agents (see Figures 5.21B and 5.23 and Tables 5.4B and 5.5). 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 stars 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 (and in Section 5.5.4), a neutral agent assigning a non-zero 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 normalized score of the Opt-2 agents is $S = 43\%$, 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 at the end of the experimental session among the five individuals of their group. In this condition, the Optimized agent (Opt-3) competing against four Mimic agents behaves as a defector (see Figures 5.21C and 5.24 and Tables 5.4C and 5.5). After ten games (like for the human subjects in our experiment), the Opt-3 agent has on average a rank of 1.7 (see Figure 5.24G). 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, $F(\mathbf{P}(t = 20), \mathbf{V})$ (see Figure 5.25 and Table 5.5). If the number of rounds was infinite, the optimal strategy for these agents (Opt-4) would be to explore the grid 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 Equation (5.1)). By using this strategy, the agents achieve a fidelity of 0.76 at round 20 (compared to 0.4 in Figure 5.5), and the fidelity would ultimately converge to 1 in the limit of an infinite number of rounds. Obviously, these Opt-4 agents

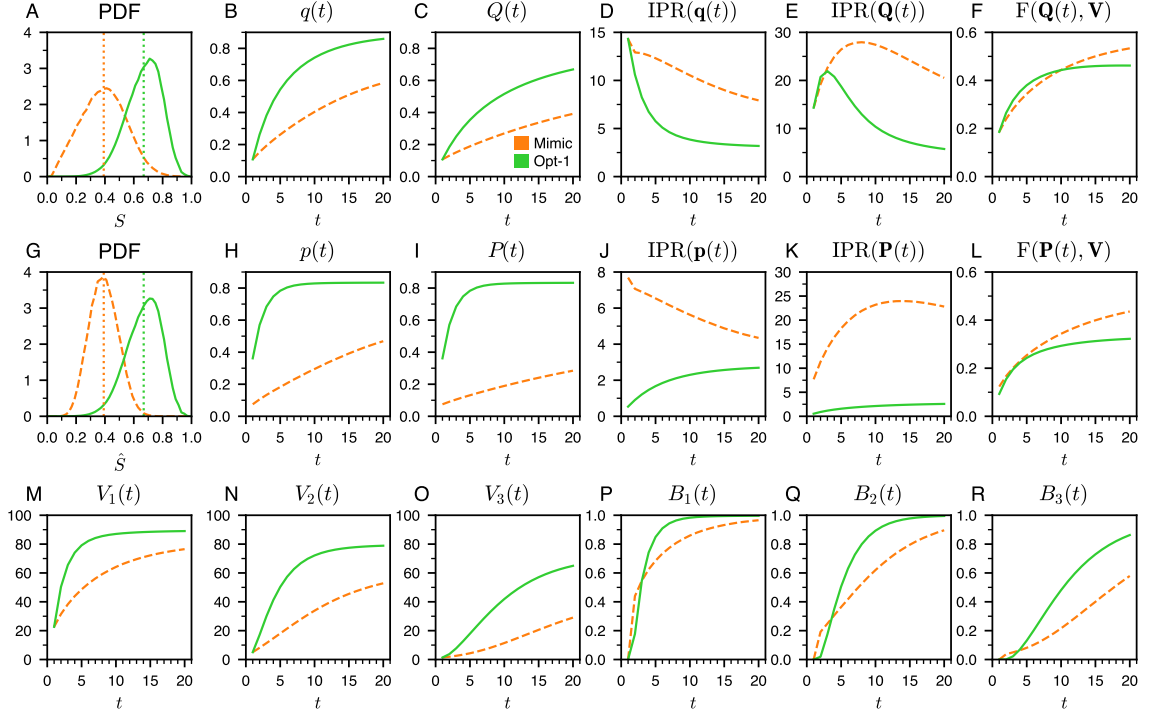


Figure 5.22: Collective performance and dynamics of collective exploration and ratings for Opt-1 agents. Simulations with five Opt-1 agents optimizing the normalized individual score S (green) compared to simulations with five Mimic agents in Rule 2 (orange). (A) Probability distribution function (PDF) of the normalized individual scores S , and (G) of the groups \hat{S} . The dotted vertical lines are the mean normalized scores. (B) Average value of the cells visited at round t , $q(t)$ and (C) up to round t , $Q(t)$. (H) Average value of the cells visited weighted by their ratings at round t , $p(t)$ and (I) up to round t , $P(t)$. (D) and (E) Inverse participation ratio of the visits, $\text{IPR}(\mathbf{q}(t))$ and $\text{IPR}(\mathbf{Q}(t))$. (J) and (K) Inverse participation ratio of the ratings, $\text{IPR}(\mathbf{p}(t))$ and $\text{IPR}(\mathbf{P}(t))$. (F) Fidelity to the cell value distribution of the distribution of visits, $F(\mathbf{Q}(t), \mathbf{V})$, and, (L) of ratings, $F(\mathbf{P}(t), \mathbf{V})$. (M–O) $V_1(t)$, $V_2(t)$, $V_3(t)$ 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 . (P–R) 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 $t > 1$.

5.5. Model Predictions

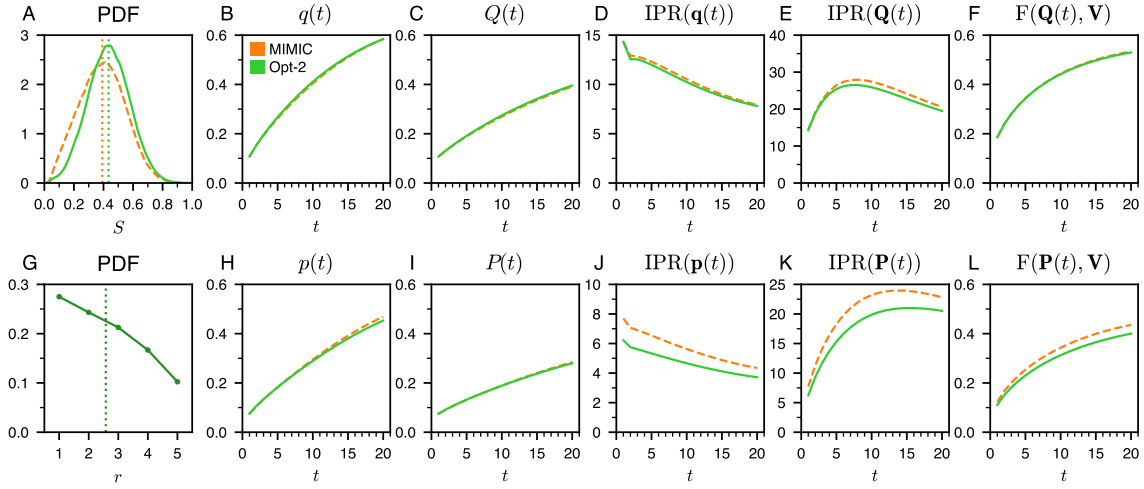


Figure 5.23: **Collective performance and dynamics of collective exploration and ratings for Opt-2 agents.** Simulations with one Opt-2 agent optimizing its score S playing with four Mimic agents in Rule 2 (green) compared to simulations with five Mimic agents in Rule 2 (orange). (A) Probability distribution function (PDF) of the normalized individual scores S . The dotted vertical lines are the mean normalized scores. (G) Probability distribution function (PDF) of the rank r of the Optimized agent after one game. The dotted vertical lines correspond to the mean rank. (B) Average value of the cells visited at round t , $q(t)$ and (C) up to round t , $Q(t)$. (H) Average value of the cells visited weighted by their ratings at round t , $p(t)$ and (I) up to round t , $P(t)$. (D) and (E) Inverse participation ratio of the visits, $\text{IPR}(\mathbf{q}(t))$ and $\text{IPR}(\mathbf{Q}(t))$. (J) and (K) Inverse participation ratio of the ratings, $\text{IPR}(\mathbf{p}(t))$ and $\text{IPR}(\mathbf{P}(t))$. (F) Fidelity to the cell value distribution of the distribution of visits, $F(\mathbf{Q}(t), \mathbf{V})$, and, (L) of ratings, $F(\mathbf{P}(t), \mathbf{V})$.

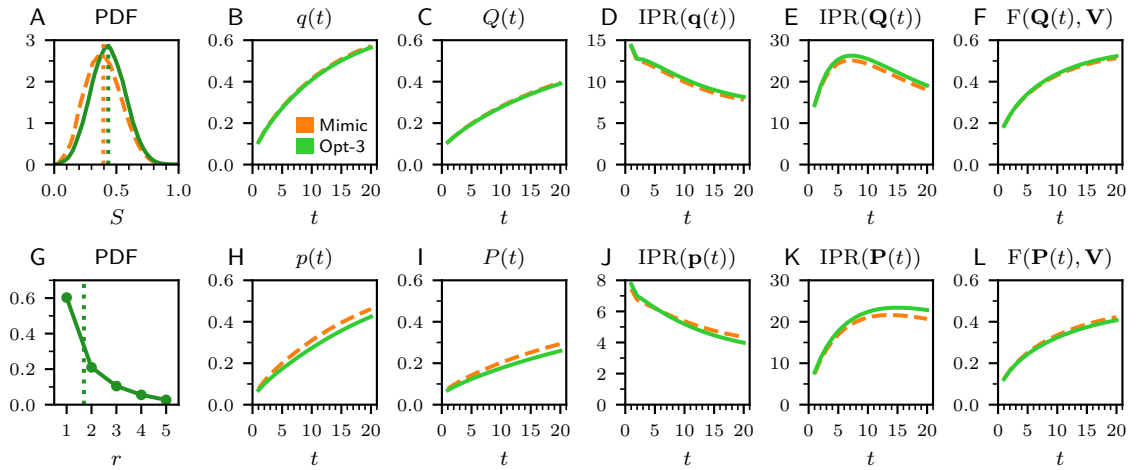


Figure 5.24: **Collective performance and dynamics of collective exploration and ratings for Opt-3 agents.** Simulations with one Opt-3 agent optimizing its rank r while playing against four Mimic agents in Rule 2 (green) compared to simulations with five Mimic agents in Rule 2 (orange). (A) Probability distribution function (PDF) of the normalized individual scores S . The dotted vertical lines are the mean normalized scores. (G) Probability distribution function (PDF) of the rank r of the Optimized agent. The dotted vertical line corresponds to the mean rank. (B) Average value of the cells visited at round t , $q(t)$ and (C) up to round t , $Q(t)$. (H) Average value of the cells visited weighted by their ratings at round t , $p(t)$ and (I) up to round t , $P(t)$. (D) and (E) Inverse participation ratio of the visits, $\text{IPR}(\mathbf{q}(t))$ and $\text{IPR}(\mathbf{Q}(t))$. (J) and (K) Inverse participation ratio of the ratings, $\text{IPR}(\mathbf{p}(t))$ and $\text{IPR}(\mathbf{P}(t))$. (F) Fidelity to the cell value distribution of the distribution of visits, $F(\mathbf{Q}(t), \mathbf{V})$, and, (L) of ratings, $F(\mathbf{P}(t), \mathbf{V})$.

5.5. Model Predictions

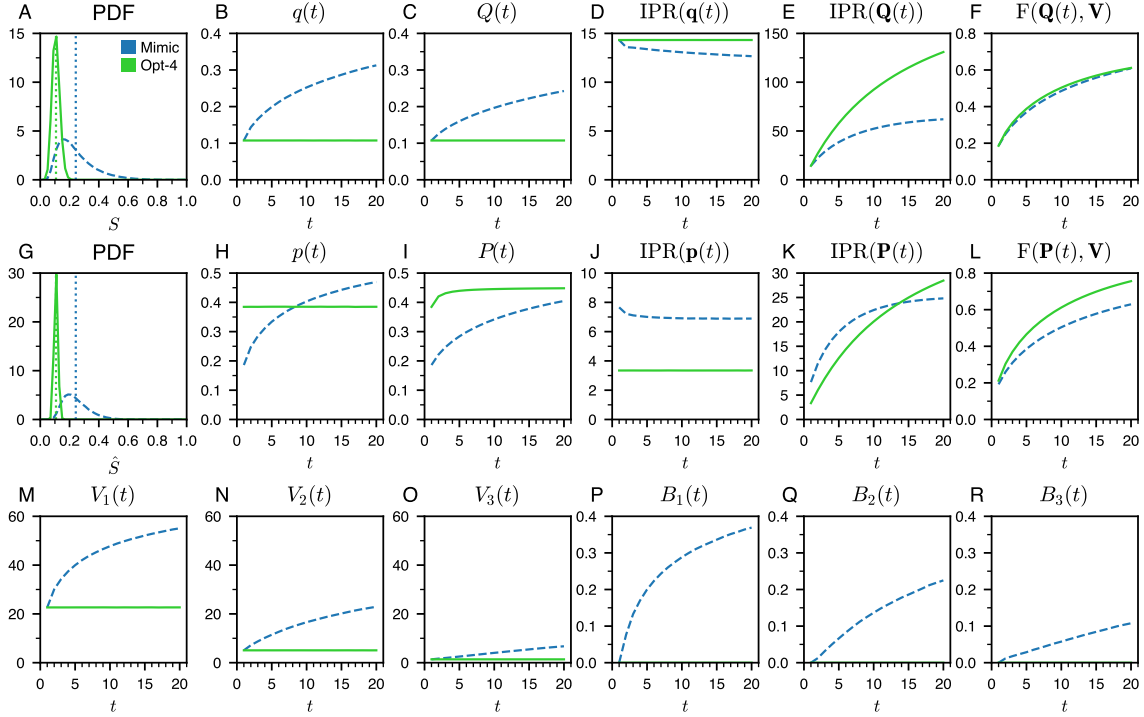


Figure 5.25: **Collective performance and dynamics of collective exploration and ratings for Opt-4 agents.** Simulations with five Opt-4 agents optimizing the fidelity of ratings with respect to cell values at the end of the experiment $F(\mathbf{P}(t = 20), \mathbf{V})$ (green) compared to simulations with five Mimic agents in Rule 1 (blue). (A) Probability distribution function (PDF) of the normalized individual scores S , and (G) of the groups \hat{S} . The dotted vertical lines are the mean normalized scores. (B) Average value of the cells visited at round t , $q(t)$ and (C) up to round t , $Q(t)$. (H) Average value of the cells visited weighted by their ratings at round t , $p(t)$ and (I) up to round t , $P(t)$. (D) and (E) Inverse participation ratio of the visits, $\text{IPR}(\mathbf{q}(t))$ and $\text{IPR}(\mathbf{Q}(t))$. (J) and (K) Inverse participation ratio of the ratings, $\text{IPR}(\mathbf{p}(t))$ and $\text{IPR}(\mathbf{P}(t))$. (F) Fidelity to the cell value distribution of the distribution of visits, $F(\mathbf{Q}(t), \mathbf{V})$, and, (L) of ratings, $F(\mathbf{P}(t), \mathbf{V})$. (M–O) $V_1(t)$, $V_2(t)$, $V_3(t)$ 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 . (P–R) 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 $t > 1$.

achieve a very mediocre mean normalized score of $S = 11\%$ compared to that of the previous Optimized agents, 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.

5.6 Solo Games

At the beginning of an experimental session, each participant would first play two games alone, so that the color of the cells on the grid would only reflect the participant's ratings. These games had two main goals: first, to ensure that participants understood the rules and how to use the web interface; second, to measure the spontaneous behavior of individuals when the only information available was the digital trace resulting from their own activity.

In solo games, players could only explore a maximum of $20 \times 3 = 60$ distinct cells among the 250 available. With such a low exploration, there is a non-negligible probability that players did not find any high-value cells during a game.

Upon examining the distributions of normalized scores in Rule 2 after one game (Figure 5.26A), we observe two distinct peaks. The first peak, around $S = 0.1$, aligns with scores expected from random exploration ($S = 0.11$), while the second peak, at $S = 0.38$, corresponds to typical game scores. Furthermore, Figure 5.26B, shows that the distribution of normalized scores after the two solo games, reveals a broader distribution with three distinct peaks. The first peak represents individuals who did not discover any high-value cells, the second peak represents those who found them once, and the third peak represents those who found them twice. This emphasizes that two repetitions of the game are insufficient to capture the full range of individual performance.

As expected, we find that a clear majority of participants behave as collaborators (with themselves) in solo games: 82% for participants playing Rule 1 and 71% for those playing Rule 2. This majority of collaborators is reflected in the mean solo rating $R(V)$ shown in Figure 5.27A, which consistently increases as a function of the cell value V . In Figures 5.27B and 5.27C, we show the behavioral profile of each individual as measured in solo games and in the ten subsequent games played in a group of five, for Rule 1 and Rule 2. For Rule 2, we did not find any significant correlation between the solo collaborators and their behavioral profile in a group of five. However, for both rule, we find that the neutrals in solo games tend to remain neutral in games in group. Interestingly enough, we also observe a few apparent solo defectors who could have been individuals adopting a different rating norm while playing alone: in a solo game, it would make sense to give a non-zero (say, 3) rating to very bad cells, to mark them to later avoid revisiting them in the next rounds while exploring the grid.

5.7 Discussion

This chapter explores the effect of a competitive versus a non-competitive condition on the way individuals use digital traces and cooperate to find the best values in a grid of

5.7. Discussion

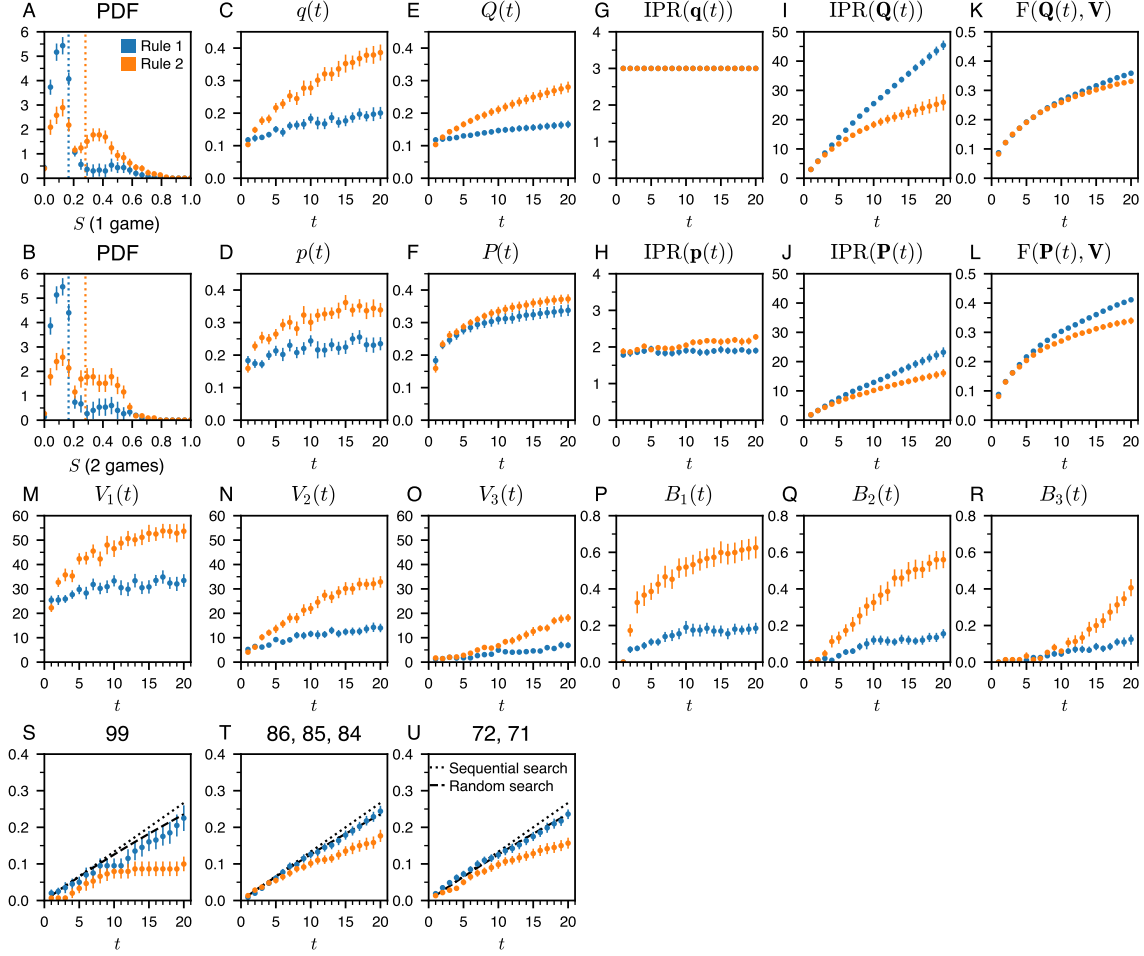


Figure 5.26: **Individual behaviors in experiments in which individuals play alone.** The non-competitive Rule 1 is in blue and the competitive Rule 2 is in orange. (A) and (B) Probability distribution function (PDF) of the normalized individual scores S , after (A) one game, and (B) two games. The dotted vertical lines are the mean normalized scores. (C) Average value of the cells visited at round t , $q(t)$ and (E) up to round t , $Q(t)$. (D) Average value of the cells visited weighted by their ratings at round t , $p(t)$ and (F) up to round t , $P(t)$. (G) and (I) Inverse participation ratio of the visits, $\text{IPR}(\mathbf{q}(t))$ and $\text{IPR}(\mathbf{Q}(t))$. (H) and (J) Inverse participation ratio of the ratings, $\text{IPR}(\mathbf{p}(t))$ and $\text{IPR}(\mathbf{P}(t))$. (K) Fidelity to the cell value distribution of the distribution of visits, $F(\mathbf{Q}(t), \mathbf{V})$, and, (L) of ratings, $F(\mathbf{P}(t), \mathbf{V})$. (M–O) $V_1(t)$, $V_2(t)$, $V_3(t)$ 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 . (P–R) 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 $t > 1$. (S) Probability to find the best cell, of value 99. (T) Probability to find one of the four cells whose values are 86 ($\times 2$), 85, or 84. (U) Probability to find one of the four cells whose values are 72 ($\times 2$) or 71 ($\times 2$).




Figure 5.27: **Mean rating and behavioral profiles in solo games.** (A) Mean rating $R(V)$ as a function of the cell value V for the experiments in which individuals play alone. The non-competitive Rule 1 is in blue and the competitive Rule 2 is in orange. (B) and (C) Change in individuals' behaviors between the single-player and five-player experiments, for (B) Rule 1 and (C) Rule 2. The x -axis represents the average slope u_1 of individuals over the two experiments in which they play alone, while the y -axis represents the average slope u_1 of individuals over the ten experiments in which they play in groups of five. The two horizontal lines at $u_{\text{def-neu}} = -0.5$ and $u_{\text{neu-col}} = 0.5$ are the delimitations between the profiles. The percentages indicate the fraction of each behavioral profile: collaborators (green), neutrals (brown), and defectors (red).

hidden numbers. Two experiments with different rules were considered. Rule 2 introduces a monetary incentive for participants to perform well, resulting in explicit competition, which is absent in Rule 1.

Our experimental results indicate 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 grid of hidden numbers. These traces constitute a form of long-term collective memory of the past actions performed by the group [Thierry et al., 1996; Baltzersen, 2022]. 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 the defectors adopt an opposite behavior. Neutral individuals constitute a sizable fraction of the group members (13% in Rule 1 and 43% 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.

5.7. Discussion

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 [Zand, 1972; Steinel and De Dreu, 2004]. This clearly occurs in Rule 2, where 82% 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 often revisiting their best-discovered cells, 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, IPR, probability of revisiting cells depending on their values. . .). 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 grid) in the choice of cells that are visited.

The analysis of the model indicates 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 to locate the best food sources in their environment [Goss et al., 1989; Beckers et al., 1990]. 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 grid, number of rounds...), and to investigate the optimal agents' strategy depending on diverse specified objectives. Our analysis indicates that the maximal score is obtained for collaborative agents (Opt-1), suggesting that intragroup collaboration should emerge from intergroup competition (see Chapter 7). Interestingly, the model also predicts that a defector behavior emerges for an agent (Opt-3) aiming at optimizing its rank, 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 obvious 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. It must be noted that despite the model's imperfection in reproducing certain observables, even the worst agreement between experimental and model results still remains within two experimental standard errors (for instance, see Figure 5.4B for Rule 1). Considering the number and diversity of observables that we have considered, 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.

Impact of Group Composition on Individual Behavior and Stigmergic Cooperation

Chapter Summary

This chapter presents the results and the modeling of nine different experimental conditions, in which each human participant competes with four social bots controlled by our model. Participants were unaware of the presence of the bots, believing that they were playing against each other in a group of five. By controlling the behavior of the bots, the setup provides a controlled environment suitable for studying human behavior and its adaptation to varying conditions. In particular, the use of different fractions of collaborator and defector bots allows for precise control of the level of cooperation within a game.

On average, the human participants can outperform the simple bots in all situations. In particular, the more collaborative the bots, the better the humans perform compared to the bots. In addition, human participants are more likely to adopt deceptive behavior when facing more collaborative groups.

A model is then constructed to mimic the behavior of human participants, using a calibration following the lines of Chapters 4 and 5. This model quantitatively reproduces the dynamics of visits and ratings, as well as the participant's performance, observed in the different experimental conditions involving social bots.

Our analysis suggests that human participants base their behavioral choices on three main natural cues that are readily available during the dozen games played per participant and that provide assessments of: *(i)* the cooperation within the group; *(ii)* the legibility and discriminative nature of the social information; and *(iii)* the efficiency of their strategy. A linear model is constructed based on these cues to predict the proportion of each behavioral profile in each situation, understanding that cues *(i)* and *(ii)* are found to be the most relevant, within a linear model.

The content of this chapter follows the lines of [Bassanetti et al. \[2024a\]](#), which will be submitted soon.

Contents

6.1 Introduction	98
6.2 Experimental Design	99
6.2.1 The experiments	99
6.2.2 The bots	100
6.3 Impact of Bots on Participant’s Performance	102
6.4 Model of the Visit and Rating Strategies	111
6.5 Model for predicting the Human Participants’ Behavioral Profiles	114
6.5.1 Cues available to human participants	114
6.5.2 Linear model for predicting individual behavioral profiles	116
6.5.3 Application of the model	117
6.5.4 Interpretation of the model	119
6.6 Discussion	120

6.1 Introduction

In Chapter 5, we demonstrated that individuals can cooperate through indirect interactions in the context of the STIGMER game. The results of our experiments also revealed that the participants exhibited distinct behavioral profiles that could be identified by their rating strategy. Collaborators assigned low ratings to cells with low values and high ratings to cells with high values. Conversely, defectors exhibit the opposite pattern. Neutrals consistently provided similar ratings regardless of the cell values. We developed a model of human behavior that implements the way individuals played the game and the different rating strategies associated with each behavioral profile.

However, the model does not explain how individuals react and adapt the way they explore and rate cells to the behaviors of the other players in the group, depending on their experience of the game. In other words, the model does not predict the composition of the group in terms of collaborators, neutrals, and defectors, and is fed the respective fractions observed experimentally. One of the main motivations of this chapter and [Bassanetti et al. \[2024a\]](#) is to gain a deeper understanding of the cues that lead a human participant to adopt a specific behavioral profile.

In the series of experiments reported in Chapter 5, there is a wide variety of behavioral profiles within each group. In the absence of competition, most group members are collaborators, whereas in the presence of competition, the groups are mainly composed of neutrals and defectors. However, not all possible combinations of the three behavioral profiles are observed in the experiments, and even in similar group compositions, there are variations among individuals with the same behavioral profiles. To address these issues, it is necessary to introduce a method for obtaining a controlled composition of the groups. This would allow for a precise investigation of their impact on individual behavior.

6.2. Experimental Design

This chapter presents and analyzes a series of nine experimental conditions in which a human participant interacts with four simple model-controlled social bots. The experimental design provides precise control over a player’s environment and ensures replicability under consistent conditions. The chapter begins with an overview of the experiments, which includes a description of the different types of bots that are used. Then, I study the impact of the level of cooperation of the bots that play with the human participants on the participant’s performance. Then, I use the model introduced in Chapter 4 to model the visit and rating strategies of human participants and better understand the mechanisms at work in the experiments. Finally, I delve deeper into the understanding of the behavioral profiles of humans. More precisely, I look at the cues available to human participants to gauge the level of cooperation among other group members and adapt their behavior accordingly. I then built a linear statistical model to predict the proportion of each behavioral profile observed in the experiments.

6.2 Experimental Design

6.2.1 The experiments

This chapter presents and compares the results of nine new experimental conditions. The experiments consist of a repetition of the STIGMER game and are played with a score that is defined as the sum of the visited values during the game. Then, at the end of the experimental session, the five participants are ranked based on their cumulative score across all games and paid accordingly. The player ranked first is paid 20€, the second is paid 15€, and the three remaining players (ranked 3–5) receive 10€ each. This incentivizes participants to have the highest score in their group, creating competition among them.

In contrast to the experiments described in Chapter 5, where five participants play in the same five-player groups, in the experiments presented here, each participant plays with four model-controlled bots. This design allows us to examine how the behavior of the bots affects the behavior of the participants in a fully controlled environment. To prevent bias in participant behavior resulting from playing with bots instead of other humans, participants are unaware that they are playing with bots and believe they are playing with one another. To achieve this, participants are instructed not to communicate with each other and are unable to view each other’s screens. Furthermore, despite the independence of the five games, participants are required to wait for each other at the end of each round before moving on to the next. This feature prevents desynchronization of the games, which could cause participants to realize that others are still playing after their game has finished. Thanks to these measures, only a very few participants suspected that something “dodgy” was happening, and none of them expressed a belief that they were playing with bots. To ensure fairness in participants’ payment, which are ranked together but do not play in the same group, all participants in the same experimental session play on the same grid, with the same shuffling of values, and against identical types of bots. As mentioned in Section 2.3, all experiments involving bots were conducted at the LPT, since deceptive experimental designs are prohibited by the TSE and most economics institutions and journals (but not in social psychology or behavioral biology).

Each of the nine experimental conditions involves a different group composition. The first five conditions test various combinations of collaborator and defector bots, ranging from a scenario with four collaborators and zero defectors (4 Col – 0 Def) to a scenario with zero collaborators and four defectors (0 Col – 4 Def). The next three experimental conditions involve players interacting with four bots always giving the same rating, regardless of the cell’s value (perfect neutrals). Three different types of bots are tested: bots always giving a rating of one star (Const-1), three stars (Const-3), and five stars (Const-5). In the final condition, players interact with four bots (Opt) that were optimized to maximize the group score (the Opt-1 agents designed and analyzed in Chapter 5).

Finally, the analysis of the present chapter also includes the experiment in which five human participants play together in the same group in the presence of competition. This experiment is the same as the “Rule 2” experiment of Chapter 5. This experiment serves as a reference situation that will be compared to the one in which one human plays with four bots.

Here is a summary of the ten experimental conditions considered in this chapter, and the corresponding number of participants who each participated in approximately twelve games:

- 1 Human vs 4 Col – 0 Def bots (10 participants)
- 1 Human vs 3 Col – 1 Def bots (15 participants)
- 1 Human vs 2 Col – 2 Def bots (15 participants)
- 1 Human vs 1 Col – 3 Def bots (15 participants)
- 1 Human vs 0 Col – 4 Def bots (10 participants)
- 1 Human vs 4 Const-1 bots (10 participants)
- 1 Human vs 4 Const-3 bots (10 participants)
- 1 Human vs 4 Const-5 bots (10 participants)
- 1 Human vs 4 Opt bots (20 participants)
- 5 Humans (14 groups of 5 participants)

6.2.2 The bots

The bots used in the experiments are controlled by the model detailed in Chapter 4. The specific parameter values that determine their visit strategy can be found in Table 6.1, while the parameters governing their rating strategy are provided in Table 6.2. Additionally, the rating strategy of the bots is visually represented in (Figure 6.1).

The collaborator and defector bots mimic the behavior of humans with collaborative and deceptive behavior in games that involve five human participants. These bots were derived from the preliminary experiments conducted in 2017, resulting in slightly different strategies compared to the Mimic agents in Rule 2 described in Chapter 5. The three constant bots Const-1, Const-3, and Const-5, have a visit strategy that is identical to that of collaborator and defector bots, and therefore, comparable to that of humans. Their rating strategy offers three variations of a neutral rating, always assigning 1, 3, or 5 stars to a visited cell. Lastly, optimized bots are designed to maximize their scores while playing in groups of five identical agents (see Opt-1 agents in Chapter 5). They explore the grid until they find high-value cells, then stop exploring and repeatedly revisit these identified

6.2. Experimental Design

	ε	α	a_1	b_1	a_2	b_2	a_3	b_3
Col, Def, Const	0.73	1.17	22.4	1.36	1.2	1.45	1.3	2.34
Opt	0.0012	0.75	35.9	2.86	30.8	2.49	27.3	2.92

Table 6.1: **Visit strategy parameters of the bots.** Values of the parameters used for the visit strategy of the bots (see Equations (4.1) and (4.2)).

	c_0	d_0	e_0	f_0	c_5	d_5	e_5	f_5
Col	0.5	0.5	7.1	-8.47	0.5	0.5	45.4	3.21
Def	0.5	0.5	30.1	1.64	0.5	0.5	3.2	-3.69
Opt	0.5	0.5	71.5	-6.17	0.5	0.5	83.6	5.29

Table 6.2: **Rating strategy parameters of the bots.** Values of the parameters used for the rating strategy of the bots (see Equation (4.4)).

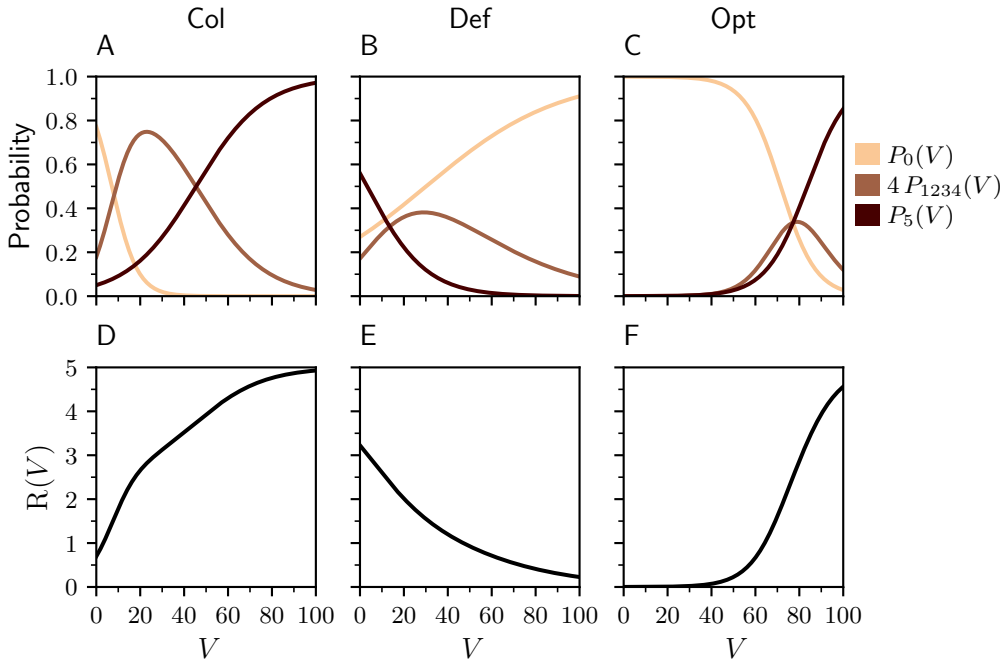


Figure 6.1: **Rating strategy of the bots.** (A–C) Probability of rating a cell with 0 stars, $P_0(V)$, 1 to 4 stars, $P_{1234}(V)$, and 5 stars, $P_5(V)$, as a function of its value V for the collaborator, defector, and optimized bots. (D–F) Mean rating $R(V)$ given to a cell of value V .

high-value cells. Their rating strategy involves rating only cells with values greater than 50 (see Figures 6.1C and 6.1F), they consequently rate only a very limited number of cells during the game.

6.3 Impact of Bots on Participant's Performance

Let us begin by examining the mean normalized score of bots and humans in each experiment (see Figure 6.2). One can observe large variability in each experiment, suggesting that the composition of the group, and consequently the behavior of the bots, have a significant impact on the overall performance of humans. Moreover, there is a strong correlation between the scores of bots and humans, indicating that the higher the bots' score, the higher the participants' score.

A closer look at the mean normalized scores for participants reveals a clear trend. The presence of more collaborator bots in a group is associated with higher mean normalized scores, with $\langle S \rangle = 0.31$ observed in the scenario in which participants played with four defector bots in the group and $\langle S \rangle = 0.56$ in the situation with four collaborator bots. In contrast, the three experiments with constant bots exhibit closely clustered mean normalized scores, averaging around 0.43. The experiment involving optimized bots secures the second-highest position with a mean normalized score of 0.48. For reference, the experiment featuring five human participants playing in the same group has a mean normalized score of 0.40, ranking among the lowest scores observed.

To better characterize the behavior of the participants, we present Figures 6.3 to 6.5, which show the same main observables as in Chapter 5 (defined in Chapter 3), but focusing on the visits and ratings of the human participant in the group. Consequently, the variables $q(t)$, $Q(t)$, $p(t)$, and $P(t)$ now represent the instantaneous and cumulative visit and rating performance of the human participant, $V_1(t)$, $V_2(t)$, and $V_3(t)$ denote the average value

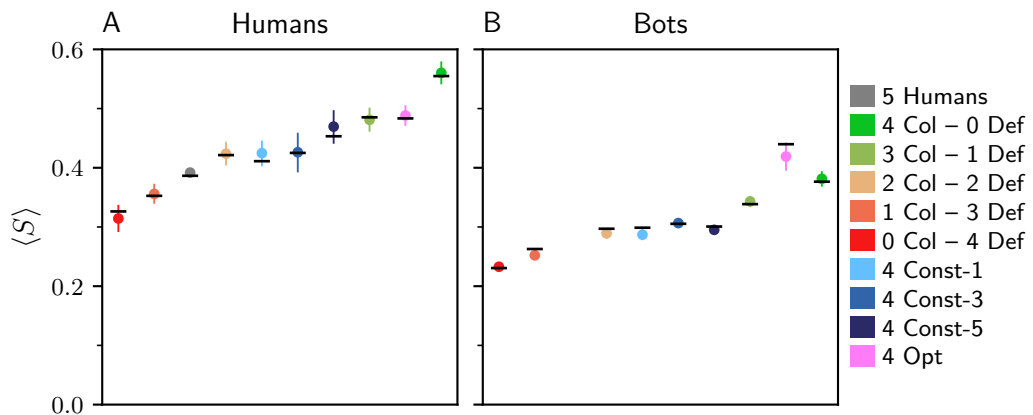


Figure 6.2: **Mean normalized score.** Mean normalized score $\langle S \rangle$ of (A) the human participants and (B) the bots in every experiment. Experiments are listed in ascending order of the mean normalized score of humans. The dots are the experimental data, and the black lines are the predictions of the model.

6.3. Impact of Bots on Participant’s Performance

of the cells visited by the human participants, and $B_1(t)$, $B_2(t)$, and $B_3(t)$ quantify the probability that a human participant revisits cells from the previous round. In addition, Figure 6.6 shows the observables for the experimental condition in which five humans played together for reference.

First, let us examine the experiments in which human participants played with collaborator and defector bots. Humans get higher normalized scores in experiments where there are more collaborator bots (Figure 6.3A). Moreover, the more collaborative is the environment, the higher the value of the cells opened (Figures 6.3B, 6.3C and 6.3F to 6.3H) and rated (Figures 6.3D and 6.3E) by the human participants. Finally, the probability of revisiting the cells visited in the previous rounds (Figures 6.3I to 6.3K) suggests that humans revisit slightly more when there are more collaborator bots.

In contrast to the diversity of performances observed in experiments with collaborator and defector bots, the three experiments involving constant bots are remarkably similar (see Figure 6.4). Therefore, this suggests that the variations in the rating of the different constant bots have a similar influence on human performance. Note that because constant bots revisit and rate the cells with the best values, they have a similar effect as collaborators, in the long run. In fact, in these scenarios, the participants’ behavior is very similar to that observed in configurations with two collaborator bots and two defector bots. The values of the observables obtained in the experiment with optimized bots are similar to those of the experiment with three collaborator bots and one defector bot (see Figure 6.5). Finally, the values of the experiment with five humans playing together are similar to those of the experiment with one collaborator bot and three defector bots (see Figure 6.6), which is consistent with the analysis of Chapter 5.

Figure 6.7 shows the probability for human participants to find the highest-value cells in experiments performed with the different kinds of bots. As for the mean normalized score, the probability of finding the highest-value cells increases with the level of collaboration within the group. The experiment with five humans stands out as having one of the lowest probabilities.

The rank of the human participants (see Figure 6.8) shows a consistent pattern. On average, humans outperform bots in all experiments. However, there is a significant difference in the mean ranks between each experiment. In situations with collaborator and defector bots, human participants tend to achieve a better ranking when there are fewer defector bots, which reduces the quality of the social information embedded in the trace. Again, the results for the constant bot experiments are almost the same, with very close rank distributions. Finally, in experiments with optimized bots, while the majority (66%) of participants secures the first position, a significant proportion (19%) finishes last. This implies that while the majority of the participants understand and benefit from the special strategy of optimized bots, a significant proportion fails to understand it, resulting in a last-place finish.

To explain the lower score of the bots compared to humans, one must note that a given bot had no option to adapt its strategy to the behavior of the other members of its group (3 bots and 1 human) and to the properties of the resulting colored grid. As will be clear hereafter, the human participants can observe some natural cues to form a strategy, for

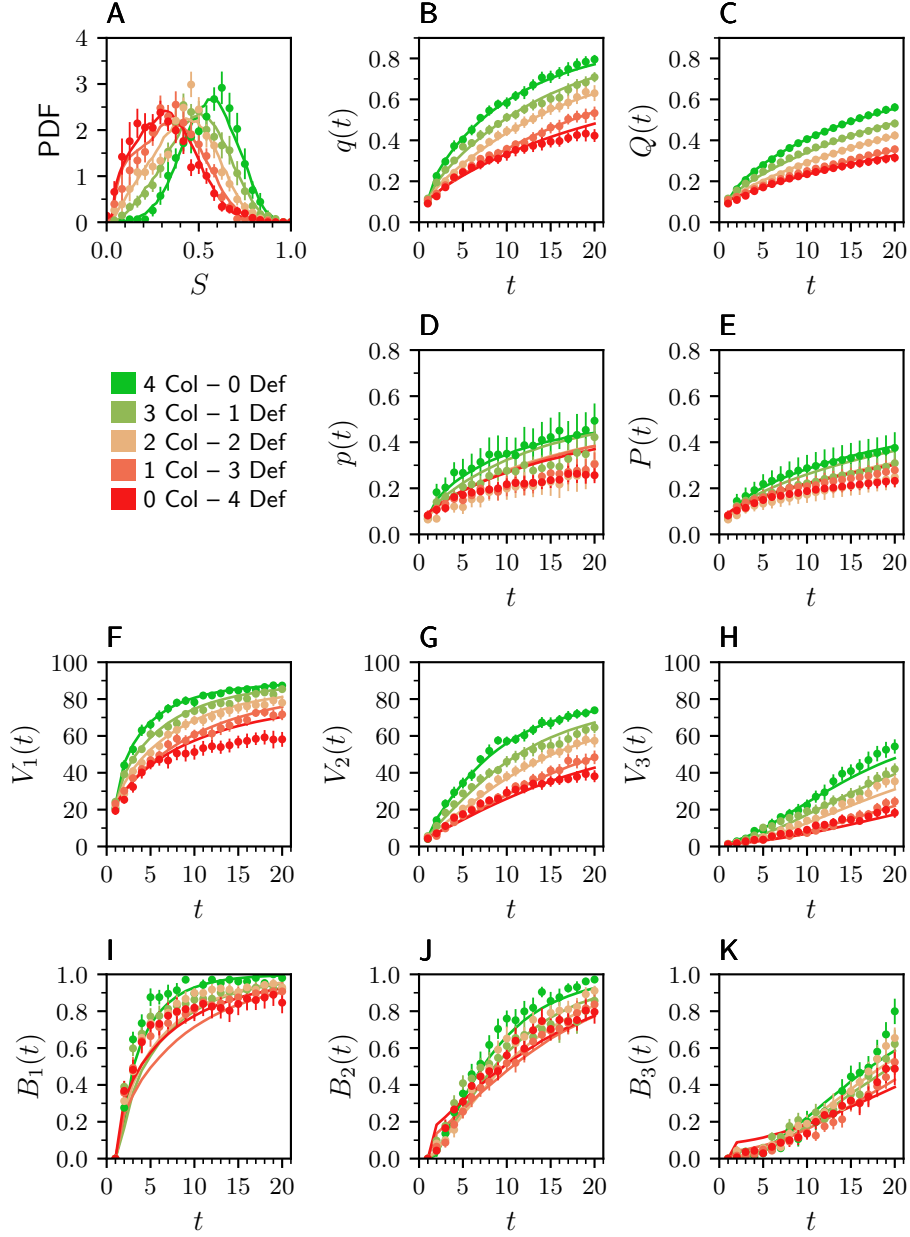


Figure 6.3: **Individual performance and behavior of human participants in experiments with collaborator and defector bots.** The dots are the experimental data, and the solid lines are the predictions of the model. (A) Probability distribution functions (PDF) of the normalized player score S of the humans (see Definition 3.9). (B) Instantaneous visit performance $q(t)$, (C) cumulative visit performance $Q(t)$, (D) instantaneous rating performance $p(t)$, and (E) cumulative visit performance $P(t)$ of the humans as a function of the round t (see Definitions 3.5 and 3.6). Value of the (F) first-best cell $V_1(t)$, (G) second-best cell $V_2(t)$, and (H) third-best cell $V_3(t)$ visited by the humans at round t (see Definition 3.11). Probability that humans revisit their (I) first-best cell $B_1(t)$, (J) second-best cell $B_2(t)$, and (K) third-best cell $B_3(t)$ visited in the previous round (see Definition 3.13).

6.3. Impact of Bots on Participant's Performance

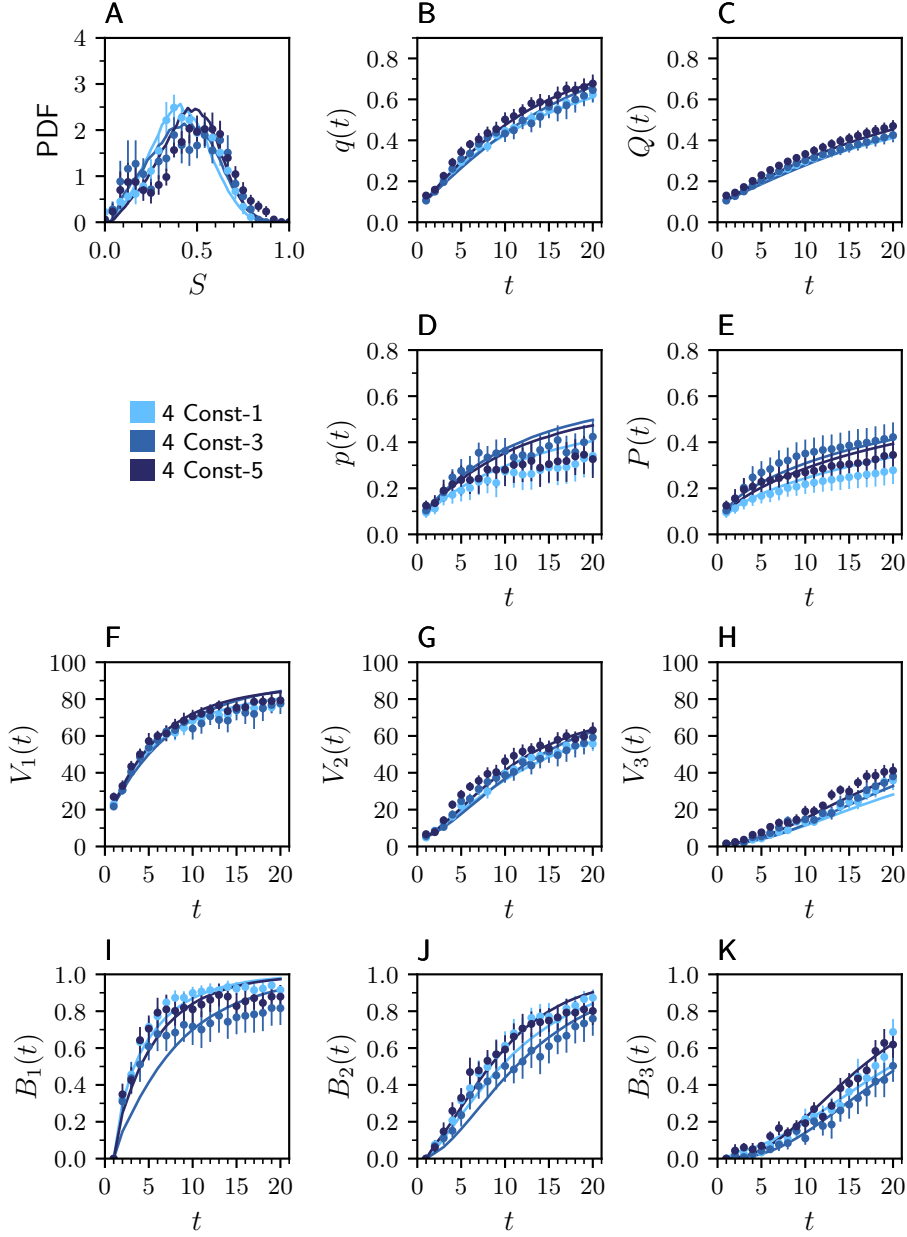


Figure 6.4: **Individual performance and behavior of human participants in experiments with constant bots.** The dots are the experimental data, and the solid lines are the predictions of the model. (A) Probability distribution functions (PDF) of the normalized player score S of the humans (see Definition 3.9). (B) Instantaneous visit performance $q(t)$, (C) cumulative visit performance $Q(t)$, (D) instantaneous rating performance $p(t)$, and (E) cumulative visit performance $P(t)$ of the humans as a function of the round t (see Definitions 3.5 and 3.6). Value of the (F) first-best cell $V_1(t)$, (G) second-best cell $V_2(t)$, and (H) third-best cell $V_3(t)$ visited by the humans at round t (see Definition 3.11). Probability that humans revisit their (I) first-best cell $B_1(t)$, (J) second-best cell $B_2(t)$, and (K) third-best cell $B_3(t)$ visited in the previous round (see Definition 3.13).

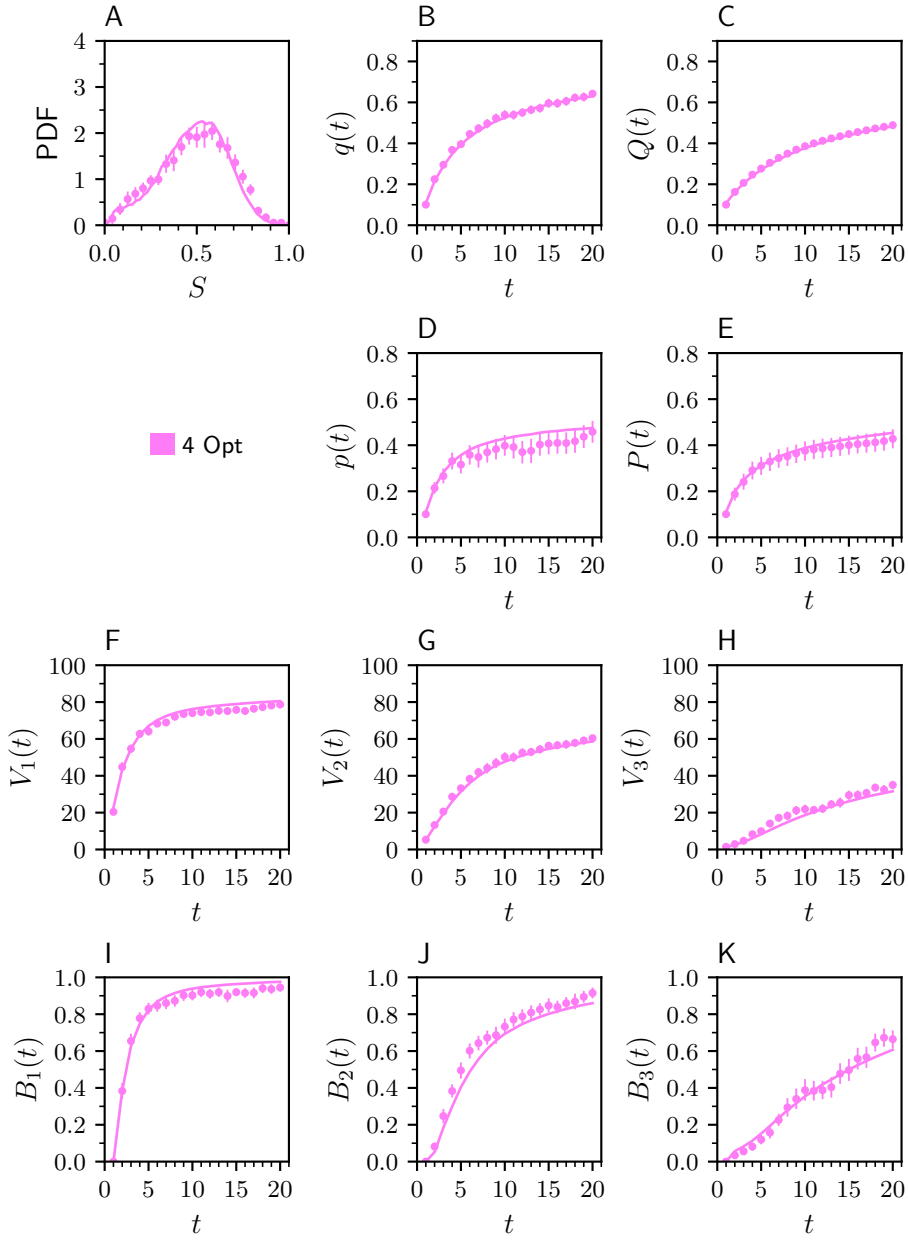


Figure 6.5: **Individual performance and behavior of human participants in experiments with optimized bots.** The dots are the experimental data, and the solid lines are the predictions of the model. (A) Probability distribution functions (PDF) of the normalized player score S of the humans (see Definition 3.9). (B) Instantaneous visit performance $q(t)$, (C) cumulative visit performance $Q(t)$, (D) instantaneous rating performance $p(t)$, and (E) cumulative rating performance $P(t)$ of the humans as a function of the round t (see Definitions 3.5 and 3.6). Value of the (F) first-best cell $V_1(t)$, (G) second-best cell $V_2(t)$, and (H) third-best cell $V_3(t)$ visited by the humans at round t (see Definition 3.11). Probability that humans revisit their (I) first-best cell $B_1(t)$, (J) second-best cell $B_2(t)$, and (K) third-best cell $B_3(t)$ visited in the previous round (see Definition 3.13).

6.3. Impact of Bots on Participant's Performance

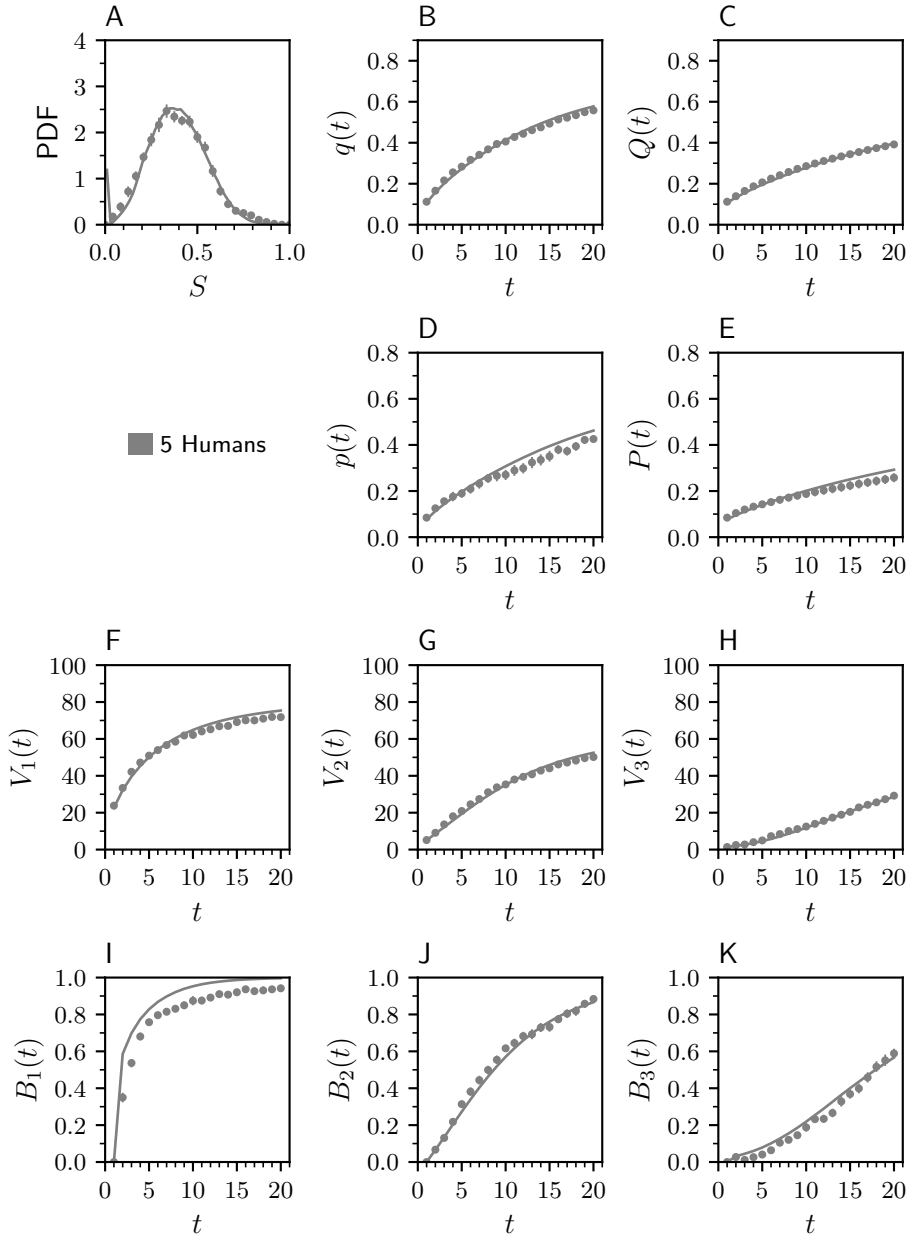


Figure 6.6: **Individual performance and behavior of human participants in experiments with groups of five humans.** The dots are the experimental data, and the solid lines are the predictions of the model. (A) Probability distribution functions (PDF) of the normalized player score S (see Definition 3.9). (B) Instantaneous visit performance $q(t)$, (C) cumulative visit performance $Q(t)$, (D) instantaneous rating performance $p(t)$, and (E) cumulative visit performance $P(t)$ as a function of the round t (see Definitions 3.5 and 3.6). Value of the (F) first-best cell $V_1(t)$, (G) second-best cell $V_2(t)$, and (H) third-best cell $V_3(t)$ visited at round t (see Definition 3.11). Probability that humans revisit their (I) first-best cell $B_1(t)$, (J) second-best cell $B_2(t)$, and (K) third-best cell $B_3(t)$ visited in the previous round (see Definition 3.13).




Figure 6.7: **Probability of finding the cells with the highest values.** Probability for a human to find (A) the best cell of value 99, (B) one of the four cells whose values are 86 ($\times 2$), 85, or 84, and (C) one of the four cells whose values are 72 ($\times 2$) or 71 ($\times 2$). Experiments are listed in ascending order of the mean probability for a human to find the best cell of value 99. The dots are the experimental data, and the black lines are the predictions of the model.

example, the observed qualitative degree of collaboration of the other members of their groups.

Section 5.4.3 introduces and defines the three different behavioral profiles that we use to classify the rating behavior of the participants: collaborators whose mean rating increases with the value of the cell, defectors whose mean rating decreases with the value, and neutral whose rating does not depend on the value of the cell. Figure 6.9 and Table 6.3 show the proportions of human participants behaving as a collaborator, neutral, and defector, in the experiments described in this chapter. For the experiments with various proportions of collaborator and defector bots, an increase in the number of collaborator bots is strongly correlated with the adoption of a more deceptive behavior by humans. This suggests that a deceptive strategy is advantageous in cooperative groups, as individuals can exploit information left by others while misleading them. However, in less collaborative groups, the effectiveness of the deceptive strategy decreases, and individuals are less often behaving as defectors. This also suggests that when faced with exceptionally low-quality information (here, due to the many defector bots), individuals turn to collaboration or a neutral behavior to leave some high-quality information for themselves.

The experiments with the optimized bots are the ones in which the human participants adopted the most collaborative behavior. And the experiments with five humans have proportions very similar to those in the situations with four collaborating bots.

6.3. Impact of Bots on Participant's Performance

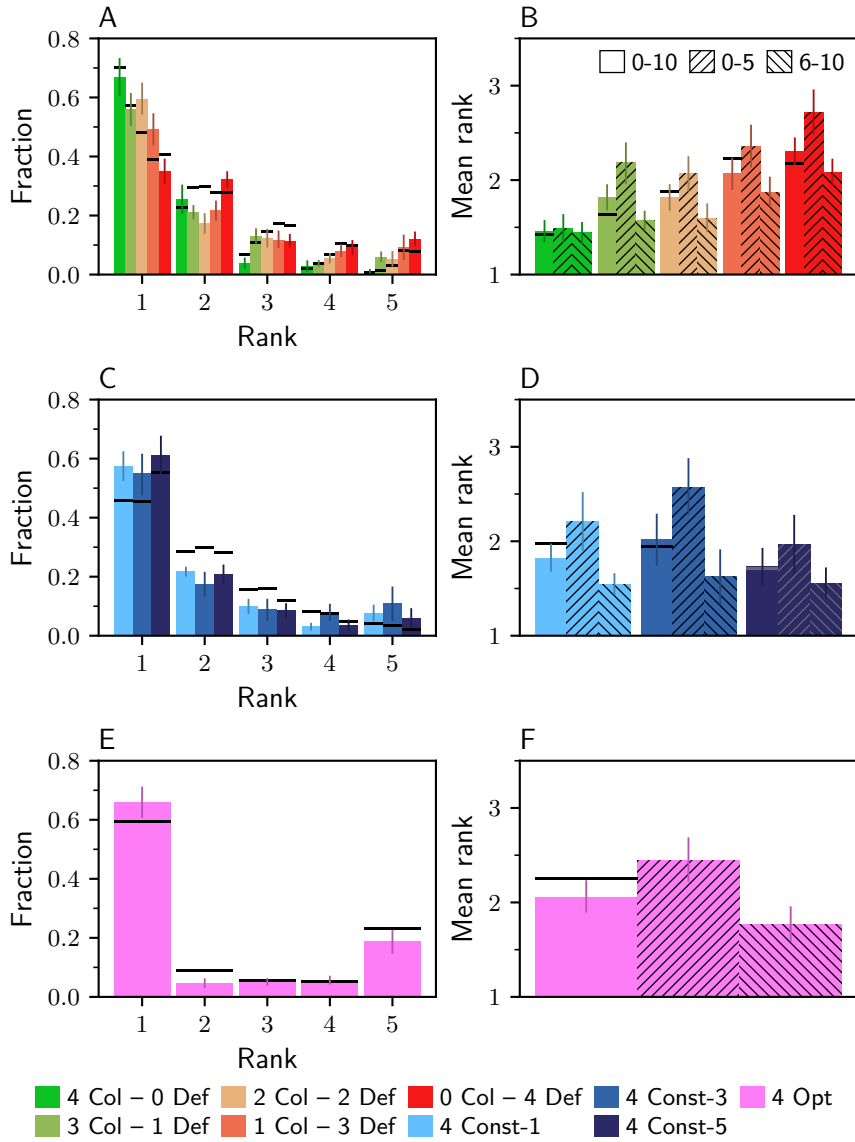


Figure 6.8: **Rank of the human participants.** (A), (C), and (E) Distribution of the rank of the human participants among the five players in the group at the end of each game. (B), (D), and (F) Mean rank of the human in the entire experimental session (left plain bar), the first five games of the session (middle hashed bar), and the last five games of the session (right hashed bar). The black lines are the predictions of the model.

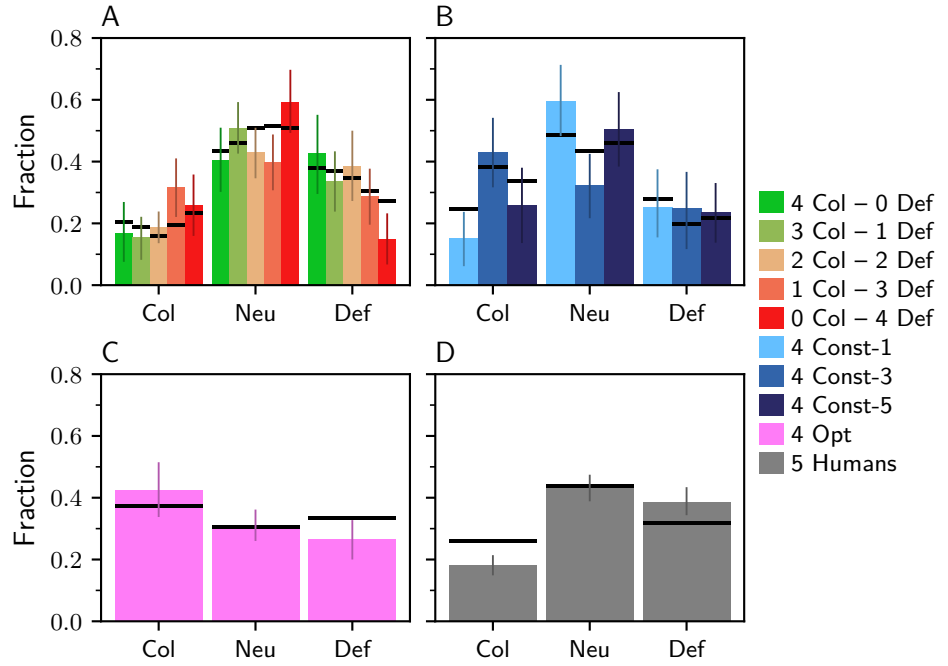


Figure 6.9: **Behavioral profile of human participants.** Fraction of collaborator, neutral, and defector for the experiments in which one human participant plays with (A) collaborator and defector bots, (B) constant bots and (C) optimized bots, and (D) for the experiment with five human participants. The black lines are the predictions of the PIR model.

	Col	Neu	Def
5 Humans	18	43	39
4 Col - 0 Def	17	40	43
3 Col - 1 Def	15	51	34
2 Col - 2 Def	19	43	38
1 Col - 3 Def	31	40	29
0 Col - 4 Def	26	59	15
4 Const-1	15	60	25
4 Const-3	43	32	25
4 Const-5	26	50	24
4 Opt	42	31	27

Table 6.3: **Behavioral profile of human participants.** Percentage of collaborative, neutral, and deceptive behaviors observed among human participants in each experiment.

6.4 Model of the Visit and Rating Strategies

Now, let us use the agent-based model described in Chapter 4 to model the behaviors of the human in each of the nine considered experimental conditions. These agents mimicking human behaviors (Mimic agents) have a strategy that is divided into two parts: the visit strategy and the rating strategy.

In each experiment, every participant’s behavioral profile—whether collaborator, neutral, or defector—is associated with a distinct rating strategy. However, we noticed that these strategies are independent of the experiments (see Figure 6.10). More precisely, a participant who adopts a deceptive strategy when playing with four collaborator bots has, on average, almost the same rating strategy as another participant with a deceptive strategy playing with four defector bots. This implies that only three rating strategies (for collaborators, neutrals, and defectors) need to be defined for all nine experiments, instead of three for each experiment.

As in Chapter 5, for simplicity, the probabilities of rating a cell with 1 to 4 stars, have been combined. Then the probabilities to rate a cell with 0 and 5 stars have been fitted using a sigmoid function (see Equation (4.4)) for collaborators and defectors, and using a linear function (see Equation (4.6)) for neutrals. The probabilities to rate a cell of value V with 1, 2, 3, or 4 stars are all equal and given by $P_{1234}(V) = (1 - P_0(V) - P_5(V))/4$. The resulting probabilities are displayed in Figure 6.11 and the values of the parameters used for the model are shown in Table 6.4. Overall, these probabilities are extremely similar to those obtained using only the experiment with five humans (see Figure 5.17).

Simulations were then run with one mimic agent playing with four bots. In each

	c_s	d_s	e_s	f_s
$s = 0$	833.7	833.6	-92.6	-4.09
$s = 5$	516.4	518.6	-1229.8	0.25

(A) Collaborator

	c''_s	f''_s
$s = 0$	0.47	0.16
$s = 5$	0.01	0.16

(B) Neutral

	c_s	d_s	e_s	f_s
$s = 0$	0.12	0.78	2.7	4.20
$s = 5$	0.37	0.36	9.9	-19.98

(C) Defector

Table 6.4: **Rating strategy parameters of Mimic agents.** Parameters values used for the rating strategy (see Equations (4.4) to (4.6)) for Mimic agents (collaborator, neutral, and defector) in the experiments with bots and in the experiment with five humans.

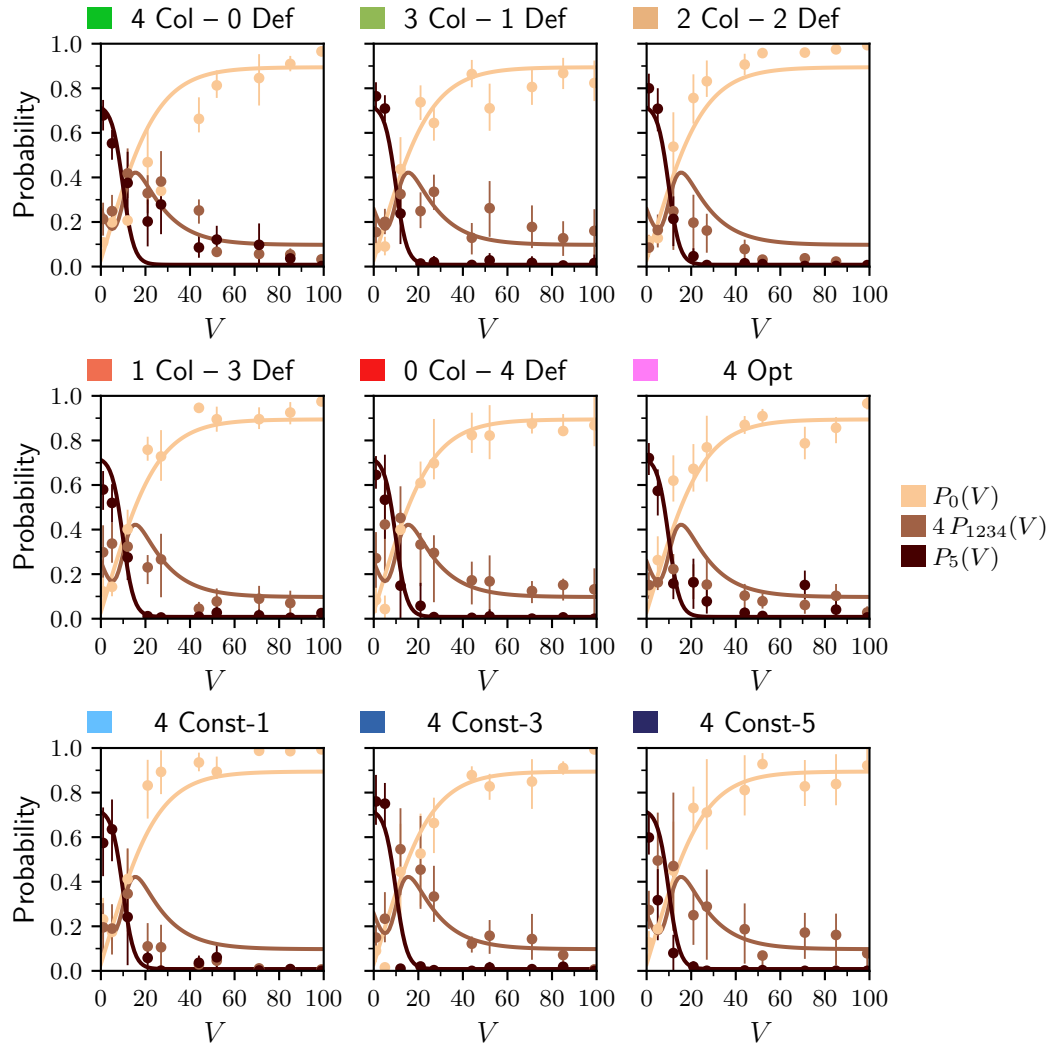


Figure 6.10: **Rating probabilities of defectors in each experimental condition.** Probability of rating a cell with 0 stars ($P_0(V)$), 5 stars ($P_5(V)$), and from 1 to 4 stars ($P_{1234}(V)$) for human participants with a defector behavior in each experiment (dots). The lines correspond to the rating strategy of a Mimic bot with deceptive behavior (see Figure 6.11C and Table 6.4C)

6.4. Model of the Visit and Rating Strategies

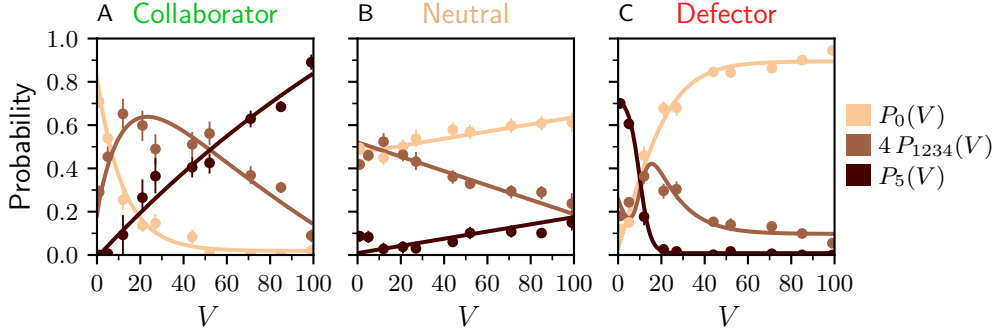


Figure 6.11: **Rating probabilities of Mimic agents.** Probability of rating a cell with 0 stars ($P_0(V)$), 5 stars ($P_5(V)$), and from 1 to 4 stars ($P_{1234}(V)$) for the collaborators, neutrals, and defectors, averaged over the nine experimental conditions. The dots are the experimental data, and the solid lines are the model.

	ε	α	a_1	b_1	a_2	b_2	a_3	b_3
5 Humans	0.69	1.32	-8.4	1.55	-4.1	2.11	-0.2	2.33
4 Col - 0 Def	0.25	1.52	19.7	2.01	18.3	1.97	20.3	1.96
3 Col - 1 Def	0.17	1.42	26.1	1.90	10.2	2.04	15.7	2.22
2 Col - 2 Def	0.23	1.12	13.0	2.00	5.4	2.10	1.5	2.23
1 Col - 3 Def	0.19	1.20	14.3	2.11	-1.5	2.09	-0.5	2.25
0 Col - 4 Def	0.58	1.04	5.8	2.16	-3.2	2.29	-2.7	2.17
4 Const	0.49	1.29	13.2	2.13	11.0	2.11	1.6	1.94
4 Opt	0.46	1.24	-2.0	1.78	4.5	2.25	-1.1	2.27

Table 6.5: **Visit strategy parameters of Mimic agents.** Parameters values used for the rating strategy (see Equations (4.1) and (4.2)) for Mimic agents in the experiments with bots and in the experiment with five humans. Note that the value of the parameter α correlates positively with the conditions leading to high-quality social information.

simulation game, the Mimic agent’s behavioral profile, or rating strategy, was randomly determined based on the observed fractions in the experiment. The visit strategy was then derived by minimizing the error between experimental and simulated observables (see Equation (4.7) and Table 6.6). The parameters defining the visit strategy are presented in Table 6.5. Note that since the observables are almost identical in the three constant bot experiments (see Figure 6.4), the visit strategy of the Mimic agents is the same in these three situations.

Looking at the values of the threshold parameters a_1 , a_2 , and a_3 (see Section 4.2.1) for the experiments in which humans played with collaborator and defector bots, we can see that the threshold at which individuals start revisiting cells from the previous round is lower when the participants are playing with many defectors. Thus, when playing with defectors, participants settle for lower values instead of continuing to search for higher values. In addition, in the nine experiments with bots, we find that the value of the parameter α defined in Section 4.2.2 correlates positively with the conditions leading to

	Δ
5 Humans	0.005
4 Col – 0 Def	0.007
3 Col – 1 Def	0.010
2 Col – 2 Def	0.021
1 Col – 3 Def	0.012
0 Col – 4 Def	0.024
4 Const-1	0.014
4 Const-3	0.021
4 Const-5	0.020
4 Opt	0.009

Table 6.6: **Error between experimental and simulated observables.** Error Δ (see Equation (4.7)) between experimental and simulated observables for every experimental condition.

high-quality social information. α is maximum when the group contains four and then two collaborator bots, and minimum when the four bots are defectors. Recall that a high value for α would correspond to a strong preference to visit the highly marked/dark cells, while a small value for α would lead to a more homogeneous selection of cells among the marked ones (see Section 4.2.2, and in particular, Equation (4.2)). Hence, we find that in the conditions where the social information is trustworthy (i.e., dark cells corresponding to higher values than light cells), the human participants consistently tend to give a larger credit to the cell colors on the grid. This also indicates that the human participants are well aware of the degree of collaboration of the four other members of their group, and can then adapt their visit and rating strategy according to this qualitative observation. Note that this also explains the better scores achieved, on average, by the participants compared to the bots. Indeed, a given bot had no option to adapt its strategy to the behavior of the other members of its group (3 bots and 1 human) and to the properties of the resulting colored grid. In fact, the present analysis (see also hereafter) paves the way to designing bots able to adapt to their environment (see Bassanetti et al. [2024a]).

Overall, looking at Figures 6.2 to 6.8, we find that the model is in good agreement with the experiment (see Table 6.6), showing that the model gives an accurate representation of the participant’s behavior.

6.5 Model for predicting the Human Participants’ Behavioral Profiles

6.5.1 Cues available to human participants

In the previous section, we built a model of human behavior in the nine experimental conditions. However, this model does not predict nor explain the distribution of the observed behavioral profiles in each condition. Therefore, it is worth investigating further

6.5. Model for predicting the Human Participants' Behavioral Profiles

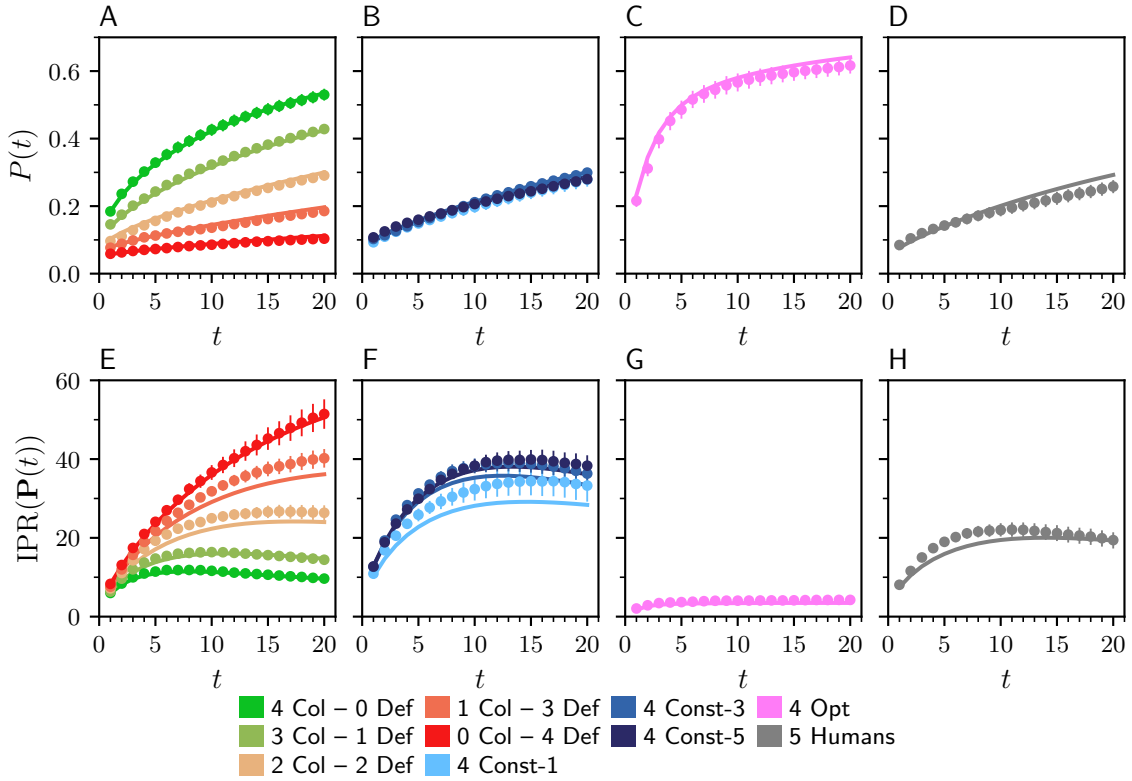


Figure 6.12: **Two cues available to human participants.** (A–D) Average value of the cells visited weighted by their ratings up to round t $P(t)$ (see Definition 3.6). Inverse participation ratio of the cumulative fraction of stars $\text{IPR}(\mathbf{P}(t))$ (see Definition 3.7). The dots represent the results for the experiments in which one human participant plays with collaborator and defector bots (A and E), constant bots (B and F), optimized bots (C and G), and for the reference experiment with five human participants (D and H). The lines correspond to the model predictions.

the factors that influence an individual's decision to engage in cooperative, neutral, or deceptive behavior in response to the actions and behaviors of the other participants resulting in a given social information (the colored grid).

Figure 6.9A shows that the more collaborator bots in the group, the more the humans choose to adopt a defector behavior. Therefore, this suggests that the presence of a high degree of collaboration in the game pushes the human to adopt deceptive behavior. In fact, the measure of the parameter α discussed in the previous section strongly indicates that the participants had a rather clear perception of the trustworthiness of the colored grid, and hence, of the qualitative degree of collaboration of the other members of their group. But, in practice, how do individuals perceive the level of cooperation of other players, and is this the only factor that matters when choosing a behavior?

Let us introduce some natural qualitative cues available to a human participant in a game to assess its environment and the properties of the social information. First, an

individual can evaluate whether the highly colored cells correspond to high- or low-value cells. This evaluation is embedded in the observable $P(t)$, which is the average value of the colored cells weighted by their respective ratings from the beginning of the game (see Definition 3.6). Figures 6.12A to 6.12D show that there is a significant variation in $P(t)$ across the different situations. In the experiment in which participants play with collaborator bots, the ratings correspond predominantly to high-value cells, while in the case in which they play with defector bots, the ratings are on low-value cells. Figure 6.12C shows in the experiment with optimized bots, the ratings are on cells with very high values. The experiment in which humans play with constant bots (Figure 6.12B) and the one in which five humans play together (Figure 6.12D) show ratings similar to the experiment in which humans play with two collaborator bots and two defector bots. Therefore, $P(t)$, which can be qualitatively evaluated by the players (especially in the later rounds) by visiting cells that have been visited and rated by others, provides a reliable indication of the level of cooperation of the other group members. However, one can see that in the experiments with optimized bots and in the one with four collaborator bots, which have the highest values of $P(t)$, the fraction of collaborators is very different. Thus, $P(t)$ is probably not the only relevant cue used by players to determine their behavior. Yet, in the next section, we will define regression models including $P(20)$ (at the end of a game) as a quantifier of the collaboration present in the group.

Another cue easily available to individuals is the complexity of the grid which corresponds to the apparent number of different cells that have been rated. Although a large number of colored cells can be beneficial if their values are high, an excessive number of colored cells muddles the social information and may also indicate that some players are defectors. The effective number of colored cells at any time is given by the Inverse Participation Ratio (IPR) of $\mathbf{P}(t)$ (see Definition 3.7). In the experiments with collaborator and defector bots, Figure 6.12E shows that there is a wide range and variation in the values of $\text{IPR}(\mathbf{P}(t))$. In the experiments with many collaborator bots, the IPR is low, and the more defector bots in the group, the higher the IPR. Figure 6.12G shows that the IPR is very low in the experiments with optimized bots because, compared to the collaborator bots, these bots only rate the cells with very high values. In the next section, we will define regression models including $\text{IPR}(\mathbf{P}(20))$ (at the end of a game) as a quantifier of the effective number of different cells that have been rated.

Finally, another natural cue accessible to human participants is their rank among the five players in the group at the end of each game (see Figure 6.8; remember that each individual was playing a dozen games during a one-hour session). This rank was explicitly displayed by the web interface at the end of each game (see Figure 2.6B). A low/good rank indicates an effective strategy, while a high/bad rank suggests room for improvement, and possibly, the need for a change in strategy.

6.5.2 Linear model for predicting individual behavioral profiles

In this section, we build a linear regression model to predict the behavioral profile of each individual in the game in the different experimental conditions, exploiting $P(20)$, $\text{IPR}(\mathbf{P}(20))$, and the rank as quantifiers of the three cues possibly driving the participants'

6.5. Model for predicting the Human Participants' Behavioral Profiles

behavior.

When working with a model, it is best to work with standardized data. We will write a standardized quantity with a hat: $\hat{X} = (X - \mu)/\sigma$ is the standardized value of the quantity X , where μ is the mean of X over the experimental data and σ is the standard deviation of X . After this standardization, \hat{X} has a zero mean and a unit standard deviation.

Let C_{exp} , N_{exp} , and D_{exp} be respectively the fractions of humans with collaborator, neutral, and defector behaviors observed in a given experimental condition, and let C_{pred} , N_{pred} , and D_{pred} be the predicted fractions. We define a feature vector of components \hat{x}_i , $i = 1, 2, \dots, f$, containing f standardized features or quantifiers expected to explain the data. In our specific case, the possible features are $P(20)$, $\text{IPR}(\mathbf{P}(20))$, and the rank, and f can go from 1 to 3.

We can now define our linear regression model,

$$\begin{cases} \hat{C}_{\text{pred}} = \sum_{i=1}^f c_i \hat{x}_i, \\ \hat{D}_{\text{pred}} = \sum_{i=1}^f d_i \hat{x}_i, \\ N_{\text{pred}} = 1 - C_{\text{pred}} - D_{\text{pred}}, \end{cases} \quad (6.1)$$

where for $i = 1, 2, \dots, f$, the c_i and d_i are regression parameters.

Then, these parameters are obtained by a fit of the model predictions to the data by minimizing the error E defined as,

$$E = \sqrt{\frac{\sum_s \left((C_{\text{exp}} - C_{\text{pred}})^2 + (N_{\text{exp}} - N_{\text{pred}})^2 + (D_{\text{exp}} - D_{\text{pred}})^2 \right)}{\sum_s \left(C_{\text{exp}}^2 + N_{\text{exp}}^2 + D_{\text{exp}}^2 \right)}}, \quad (6.2)$$

where

$$\begin{cases} C_{\text{pred}} = \mu_C + \sigma_C \sum_{i=1}^f c_i \hat{x}_i, \\ D_{\text{pred}} = \mu_D + \sigma_D \sum_{i=1}^f d_i \hat{x}_i, \\ N_{\text{pred}} = 1 - C_{\text{pred}} - D_{\text{pred}}. \end{cases} \quad (6.3)$$

Note that due to the symmetric form in C , D , N of the error, linear regressions on the observables (C and D) or (C and N) or (N and D) would result in the same predictor.

In our context, we have ten different and independent experimental conditions, and two independent variables (C and D , while $N = 1 - C - D$), representing twenty independent measurements to be explained by the linear regression model. The number of unknown parameters is $2 \times f$, where f can go from 1 to 3, depending on the number of cues used as features, among $P(20)$, $\text{IPR}(\mathbf{P}(20))$, and the rank.

6.5.3 Application of the model

Let us now apply the regression linear model to the ten experimental conditions studied in this chapter. Only using one feature ($P(20)$, or $\text{IPR}(\mathbf{P}(20))$, or the rank) in the linear

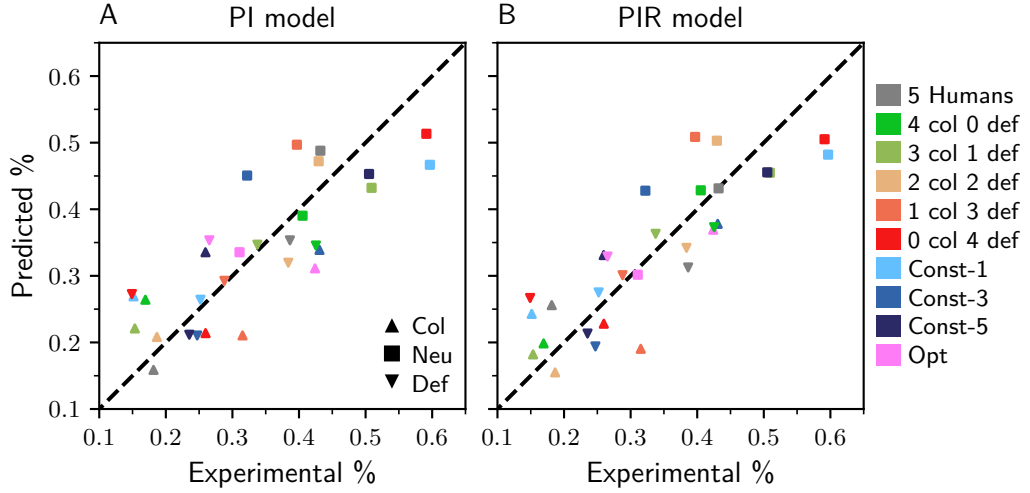


Figure 6.13: **Performance of the PI and PIR models.** Scatter plot of the predicted fractions of collaborator, neutral, and defectors as a function of the corresponding fractions observed in each experiment for (A) PI model using $P(20)$ and $\text{IPR}(\mathbf{P}(20))$ as features and (B) the PIR model using $P(20)$, $\text{IPR}(\mathbf{P}(20))$, and the rank. The dotted diagonal line represents perfect predictions.

	$P(20)$		$\text{IPR}(\mathbf{P}(20))$		Rank		E
	c	d	c	d	c	d	
PI model	1.32	-1.09	1.31	-1.68	-	-	0.204
PIR model	2.79	-2.06	2.53	-2.49	0.88	-0.59	0.159

Table 6.7: **Parameters values of the model for predicting individual behavioral profiles.** Values of the parameters and error E after optimization of the PI and the PIR models (see Equations (6.2) and (6.3)).

regression does not result in a good predictor for the fraction of the three behavioral profiles ($E = 0.239$ for the model with $P(20)$, for $E = 0.232$ the model with $\text{IPR}(\mathbf{P}(20))$, and for $E = 0.259$ the model with the rank). However, such a simple single-feature linear model already reveals that $P(20)$ and $\text{IPR}(\mathbf{P}(20))$ exhibit the best (and similar) correlation with the data, while the rank alone is certainly a less relevant feature.

We now direct our attention to a linear model involving the two main cues/features $P(20)$ and $\text{IPR}(\mathbf{P}(20))$. This model will be referred to as the “PI model”. After solving the linear regression problem to obtain the value of the four parameters, we obtain an error of $E = 0.204$ (see Figure 6.13A). Table 6.7 shows that for collaborators, the parameters for both features are almost identical and hence have the same importance. However, for defectors, $\text{IPR}(\mathbf{P}(20))$ is more correlated to the data than $P(20)$.

We now address the linear model including the three cues/features, $P(20)$, $\text{IPR}(\mathbf{P}(20))$, and the rank. This model will be referred to as the “PIR model”. After fitting the data, to obtain the value of the six parameters of the model, we obtain an improved error of $E = 0.159$ (see Figure 6.13B), but to the price of two additional fitting parameters for

6.5. Model for predicting the Human Participants' Behavioral Profiles

twenty independent data points. The inspection of the value of the six parameters (see Table 6.7) confirms that the rank is a less important feature than $P(20)$, $\text{IPR}(\mathbf{P}(20))$, as inferred from the single-feature regressions. Yet, the rank leads to significant regression parameters leading to a more accurate PIR model.

Note that in the two models described above, the quantities $P(t)$ and $\text{IPR}(\mathbf{P}(t))$ are considered at the end of the game, i.e., at the final round $t = 20$. However, by considering the round in the middle of the game ($t = 10$), the model gives similar errors ($E = 0.211$ for the model with $P(10)$ and $\text{IPR}(\mathbf{P}(10))$ and $E = 0.186$ for the model with $P(10)$, $\text{IPR}(\mathbf{P}(10))$ and the rank).

Furthermore, we conducted tests with alternative linear models incorporating other cues/features (for instance, involving the fidelity F or other qualitative markers for collaboration). However, the three cues/features, $P(20)$, $\text{IPR}(\mathbf{P}(20))$, and the rank, emerged as the most effective in terms of explanatory power.

6.5.4 Interpretation of the model

Let us have a closer look at the actual equations of the fractions of collaborators, neutrals, and defectors in the two models.

For the PI model, the fractions are given by:

$$\begin{cases} C_{\text{pred}} = \mu_C + 0.13 \hat{P} + 0.13 \hat{I}, \\ N_{\text{pred}} = \mu_N - 0.04 \hat{P} + 0.00 \hat{I}, \\ D_{\text{pred}} = \mu_D - 0.09 \hat{P} - 0.13 \hat{I}, \end{cases} \quad (6.4)$$

where \hat{P} and \hat{I} are the standardized values of $P(20)$ and $\text{IPR}(\mathbf{P}(20))$, respectively, and $\mu_C = 0.25$, $\mu_N = 0.45$, and $\mu_D = 0.30$ are the mean fractions of humans with collaborator, neutral, and defector behaviors observed in all experiments. Equation (6.4) and Figure 6.14 show that in the PI model, a larger $P(20)$ results in an increase in the fraction of collaborators and a decrease in both neutrals and defectors. The parameter $\text{IPR}(\mathbf{P}(20))$ has no discernible impact on neutrals, and as its value increases, there are fewer defectors and more collaborators.

For the PIR model, the fractions are given by:

$$\begin{cases} C_{\text{pred}} = \mu_C + 0.28 \hat{P} + 0.25 \hat{I} + 0.09 \hat{R}, \\ N_{\text{pred}} = \mu_N - 0.11 \hat{P} - 0.05 \hat{I} - 0.04 \hat{R}, \\ D_{\text{pred}} = \mu_D - 0.17 \hat{P} - 0.20 \hat{I} - 0.05 \hat{R}, \end{cases} \quad (6.5)$$

where \hat{P} , \hat{I} and \hat{R} are the standardized values of $P(20)$, $\text{IPR}(\mathbf{P}(20))$, and the rank, respectively, and $\mu_C = 0.25$, $\mu_N = 0.45$, and $\mu_D = 0.30$ are the mean fractions of humans with collaborator, neutral, and defector behaviors observed in all experiments.

In the PIR model (see Equation (6.5)), as in the PI model, an increase in the parameter $P(20)$ results in more collaborators and fewer neutrals and defectors. Additionally, an increase in the parameter $\text{IPR}(\mathbf{P}(20))$ results in an increase in the number of collaborators, a decrease in the number of defectors, and a slight decrease in the number of

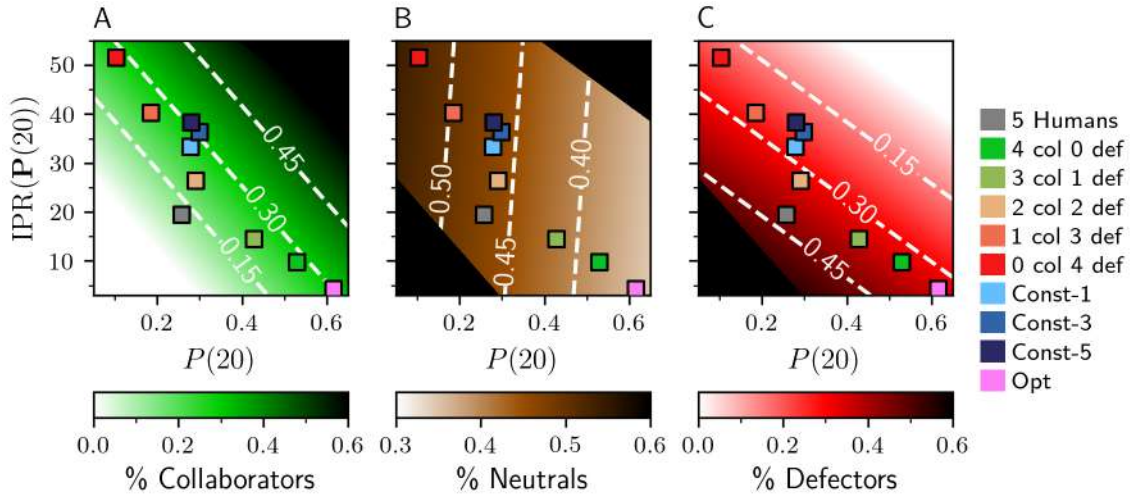


Figure 6.14: **Behavioral profiles given by the PI model.** The gradients correspond to the fractions of (A) collaborators, (B) neutrals, and (C) defectors given by the PI model as a function of $P(20)$ and $\text{IPR}(\mathbf{P}(20))$. The black and white regions are areas in which the PI model is undefined because one of the fractions is negative. Symbols indicate the position of the experimental data on the $P(20) - \text{IPR}(\mathbf{P}(20))$ plane.

neutrals. Finally, as the rank decreases, indicating lower performance, the proportion of collaborators increases, while the number of neutrals and defectors decreases.

6.6 Discussion

This chapter reports and analyzes experiments in which human participants interact with four model-controlled social bots while remaining unaware that they were playing against bots. This setup provides a controlled framework suitable for studying human behavior and its adaptation under varying conditions. The use of model-controlled bots with collaborative and deceptive behaviors enables precise control over the level of cooperation within a game, facilitating the study of various collaborative contexts.

Regardless of the collaboration level of the group, human participants were able to outperform the bots. As the level of cooperation of the bots increases, the score of human participants also increases, and they are more likely to find the highest-value cells on the grid. These results highlight the influence of group composition on cooperative dynamics and individual performance. The human participants achieved this better performance by adapting their visit and rating strategy to the behavior of the four other members of their group (bots) affecting the shared social information (the colored cells on the grid). This option to adapt to the encountered environment was not available to the bots, with their fixed visit and rating strategy during all games, and they were penalized by this asymmetry.

Furthermore, we discovered a correlation between the distribution of behavioral profiles and the level of cooperation in the game. Specifically, we observed that as the number

6.6. Discussion

of collaborator bots increased, there was a greater likelihood of human players exhibiting defector-like behavior. It appears that some human participants exploit the collaborative nature of the game to manipulate the trace and lower the performance of others.

The analysis revealed that participants based their behavioral choices on three main natural cues available during the game: *(i)* an assessment of whether the colored cells correspond to high- or low-value cells (group cooperation); *(ii)* an assessment of the number of different cells that have been rated (legibility and discriminative nature of the social information); and *(iii)* an assessment of their performance (i.e., rank) relative to other players in the game (the efficiency of their strategy).

A typical game characterized by a high level of cooperation shows distinct features: the cells that are rated by the players are predominantly those with high values, and the number of colored cells remains relatively small. Conversely, in situations characterized by a lack of cooperation, a large number of cells are rated, but these cells tend to have low or intermediate values.

To gain a deeper insight into how individuals use these cues to make decisions, we constructed a linear model, incorporating the three cues described above, to predict the proportion of each behavioral profile. This simple model yielded promising results and offered a meaningful interpretation of human behavior observed in the different controlled conditions. In particular, we found that among the three cues, the values of the colored cells and the number of colored cells are more influential than the ranking on the participants' behavior.

The linear model indicates that an increase in the average values of the highly colored cells or an increase in the number of different cells rated correlates with an increase in the number of collaborators and a decrease in the number of defectors. Furthermore, the model shows that good (low) ranks are more favorable for neutral or deceptive behavior, whereas bad (high) ranks are correlated with an increased likelihood of collaborative behavior. This suggests that participants may strategically switch from collaborator to neutral or defector to improve their ranking.

One missing element in the experiments presented so far is a notion of punishment for a deceptive behavior. In the real world, individuals who choose to deceive often face repercussions, either directly or indirectly. To incorporate such repercussions into the STIGMER game, a first option is to introduce a penalty for a deceptive rating or a bonus for a faithful and cooperative rating (the latter option was tested in preliminary experiments mentioned in Section 2.1.4).

Another option consists of introducing intergroup competition alongside intragroup competition. In this context, individuals who use a deceptive rating strategy would undermine the overall performance of their group, resulting in negative consequences for themselves (and the other members of the group). This could even lead to an overall decrease in collaboration and an increase in defection in the group, in a negative feedback loop, further undermining the performance of the group. These questions are, among others, addressed in Chapter 7, where we present and compare nine experimental conditions, including one that involves pure intergroup competition, leading to overwhelming collaboration within the groups.

Impact of Intragroup and Intergroup Competition on Stigmergic Cooperation

Chapter Summary

This chapter presents the results of seven experimental conditions with different payment schemes that introduce diverse levels of intragroup and intergroup competition between two groups of five participants. This setup allows for the analysis of the impact of intragroup and intergroup competition on human behavior in the framework of the STIGMER game.

Our qualitative analysis shows that in the absence of competition, individuals tend to cooperate, especially when rewards are uniform. However, when payment disparities based on performance are introduced, a less cooperative environment is created, even when competition is effectively absent. Additionally, payment schemes that involve intergroup competition tend to promote a cooperative behavior within a group, while those with intragroup competition tend to foster a deceptive behavior.

To quantitatively characterize the seven payment schemes, three key markers or features are introduced: *(i)* the standard deviation of rewards within groups, which we relate to an assessment of the intragroup competition; *(ii)* the absolute difference in mean reward between the two groups, which assesses the direct intergroup competition; and *(iii)* the mean difference between the five highest and five lowest rewards, which quantifies the global competition among the ten participants. We introduce linear regression models exploiting these three features to predict the proportion of each behavioral profile in the seven experimental conditions. We find that *(i)* is the feature with by far the best predictive power, already being able to differentiate the seven conditions according to their intragroup and intergroup competition. Indeed, the linear model based on the sole feature *(i)* already gives good results. Ultimately, the second most relevant feature is *(iii)*, and the model based on *(i)* and *(iii)* leads to an excellent prediction of the fraction of the three behavior profiles. Finally, we find that feature *(ii)* exhibits only a weak correlation with the data.

The content of this chapter follows the lines of [Bassanetti et al. \[2024b\]](#) (in preparation).

Contents

7.1 Introduction	124
7.2 Experimental Design	124
7.3 Performance and Behavioral Profiles of the Participants	126
7.4 Model for Predicting Individual Behavioral Profiles	130
7.4.1 Three quantities characterizing the payment schemes	130
7.4.2 Application of the model	132
7.5 Discussion	134

7.1 Introduction

So far, the analysis has focused primarily on experiments either without competition (Rule 1 in Chapter 5) or with intragroup competition (Rule 2 in Chapter 5, and all experiments with bots in Chapter 6). However, it is important to recognize that natural scenarios often involve multiple groups competing within the same environment, leading to a complex hierarchy of multilevel competition.

Multiscale models of the evolution of cooperation predict that intragroup cooperation can be favored when intergroup competition is more intense than intragroup competition [Wilson, 2015; Cooney, 2019]. To test this hypothesis, this chapter presents a series of experiments investigating the effects of various payment schemes on intragroup and intergroup competition, as well as their impact on collective performance and individual behavior.

The chapter begins with a detailed description of the experiments, followed by an analysis of the groups' collective performance and of the individual behavior of participants in different competitive conditions. Finally, a simple linear regression model is presented to interpret the participants' behavior in each specific context.

The content of the present chapter follows the lines of Bassanetti et al. [2024b] (in preparation).

7.2 Experimental Design

In this chapter, we investigate the impact of seven competition schemes that combine different levels of intragroup and intergroup competition. Each of these competition schemes makes use of different payment methods to compensate the participants at the end of an experimental session. By manipulating the payment structure, it is possible to create different types and levels of competition, both between participants of the same group (intragroup competition) and between different groups of participants (intergroup competition). Even within the same type of competition, the payment scheme can be fine-tuned to provide slightly different incentives to the players.

7.2. Experimental Design

Each experimental session involves ten participants. First, each participant plays alone two games of twenty rounds, to ensure their correct comprehension of the game and to characterize their spontaneous behavior when the only information available is the digital traces resulting solely from their own activity (see the analysis of solo games in Section 5.6). The participants are then randomly divided into two separate groups of five players each and play typically ten games of twenty rounds in their respective groups. At the end of each experimental session, the players are paid according to a specific payment structure, which depends on the experimental condition. Between seven and ten one-hour experimental sessions were conducted for each payment structure.

The first two payment structures do not create any explicit competition among participants, but correspond to different incentives:

- **No competition + equal reward:** In this condition, there is no score or ranking between the participants. Each participant received a fixed reward of 10€. Here, individuals are free to do what they want, as their behavior does not affect their reward. This experiment is the same as the one referred to as “Rule 1” and thoroughly analyzed in Chapter 5. We conducted 10 sessions of 2 groups of this experimental condition.
- **No competition + performance reward:** In this condition, there is a score, but participants are not ranked, and are told that they will be paid strictly proportionally to their individual score. The participants were also told that if they exactly achieved the average score of a game as we had measured so far in all our experiments, their reward would be close to 11€. This payment scheme encourages individuals to achieve the highest possible score, which is expected to lead to cooperation among participants identifying (with their ratings) the cells with high values. In principle, in this experiment, the participants should not resist broadcasting their best-discovered cells to the other members of their group, also hoping to elicit the same behavior from the other players. Yet, despite the total absence of ranking, one cannot avoid observing (as we did) certain participants trying to beat the other members of their group by adopting a deceptive behavior, although this could only have a negative impact on their score. We conducted 8 sessions of 2 groups of this experimental condition.

Note that in this experiment, unbeknownst to the participants, we normalized the payments to ensure that the total cost of a one-hour session would remain constant ($10 \times 11\text{€} = 110\text{€}$) so that we could anticipate the total cost of this experimental condition. Hence, in the formula to compute the reward of a player as a function of their score, $\text{reward} = 11 \times \text{score} / \langle \text{score} \rangle$, the denominator was not the average score over all experiments so far (as told to the participants in the instructions they received before playing), but the average score of the 10 players in the considered session. Therefore, the reward of a player was not truly proportional to their score but also depended on the performance of the nine other players of the session. However, since the participants were unaware of this, there was no bias, and this experimental condition was truly non-competitive in essence, albeit the players deciding otherwise.

The second type of payment structure relies solely on intragroup competition. There is competition between the members of a group, but no competition between the two groups. In each group, the five players are ranked based on their scores and paid accordingly:

- **Intragroup competition only:** In each group, the participant ranked first received 20€, the participant ranked second received 15€, and each of the three remaining participants (ranked third to fifth) received 10€. This experiment is the same as the one denoted “Rule 2” in Chapter 5 and “5 Humans” in Chapter 6. We conducted 7 sessions of 2 groups of this experimental condition.

The third type of payment structure solely involves intergroup competition. There is no explicit or implicit competition between the five members of a group, but there is an explicit competition between the two groups. This condition is expected to elicit a strong intragroup cooperation that we aimed at quantifying. The two groups are ranked based on the average score of their five members:

- **Intergroup competition only:** Each participant of the group with the highest score (average score of its members) is paid 20€, while each member of the other group is paid 10€ each. We conducted 9 sessions of 2 groups of this experimental condition.

The fourth type of payment structure combines both intragroup and intergroup competition. At the end of the experimental session, the ten players (five in each group) are ranked together based on their individual scores and rewarded accordingly. Although there is both intragroup and intergroup competition in the three experiments described below, the variations in the payment schemes change the balance between intragroup and intergroup competition in each experimental condition:

- **Inter + intragroup competition (1):** The participants ranked first through fifth are paid 15€ each, while the participants ranked sixth through tenth each receives 8€. We conducted 10 sessions of 2 groups of this experimental condition.
- **Inter + intragroup competition (2):** The participants are paid linearly according to their rank. The participant ranked first is paid 15€, the one ranked second 14€, the one ranked third 13€, and so on, down to the last participant, who is paid 6€. We conducted 10 sessions of 2 groups of this experimental condition.
- **Inter + intragroup competition (3):** The participant ranked first is paid 20€, the ones ranked second and third are paid 15€, and the remaining seven players are paid 10€ each. We conducted 8 sessions of 2 groups of this experimental condition. We conducted 8 sessions of 2 groups of this experimental condition.

7.3 Performance and Behavioral Profiles of the Participants

Let us first consider the performance of the individuals in each experimental condition. Figure 7.1 shows the mean normalized score (see Definition 3.10) and the probability of

7.3. Performance and Behavioral Profiles of the Participants




Figure 7.1: **Performance of the participants.** (A) Mean normalized score $\langle S \rangle$ of the participants in the seven experimental conditions. (B) Probability to find one of the nine cells with the highest value at round $t = 20$ in the seven experimental conditions. The black dashed and dotted lines correspond to the expected probabilities of two different visit strategies: cells chosen randomly (see Equation (3.19)), and cells chosen sequentially (see Equation (3.20)), respectively.

finding the highest-value cells (see Definition 3.14). In the condition with no competition + equal rewards (see the detailed analysis of this experimental condition in Chapter 5), participants behave in an entirely different way than in the other experiments. The mean normalized score is much lower than that obtained in the other experiments, and the probability of finding the highest-value cells is much higher. This is explained by the absence of incentives to get a high score, which leads individuals to explore the grid rather than revisiting high-value cells discovered in previous rounds. This experiment is notably the only one in which the probability of finding the high-value cells is higher than in the random and sequential visit strategies (see Equations (3.19) and (3.20)).

In the other six experiments, the mean normalized score and the probability of finding the high-value cells are ranked in the same order. The next experiment with the highest performances is the intergroup competition experiment, followed by the No comp. + performance reward experiment. However, even though these two experiments should motivate participants to cooperate, this motivation appears to be more effective in the presence of intergroup competition than for the experiment where the participants are paid proportionally to the score. As mentioned in the previous section introducing the seven experimental conditions, some participants in the No comp. + performance reward experiment were motivated by the notion of a score to perform better than the other members of their group, instead of simply focusing on the true objective, which is to maximize their score, irrespectively of the score of others. Actually, in informal discussions at the end of the No comp. + performance reward sessions, several participants admitted their personal competitive approach to this experimental condition. Moreover, in both of these experiments, the mean normalized score is much lower than that of the Opt-1 bots ($\langle S \rangle = 0.67$; see Section 5.5.5), which were optimized to maximize the mean normalized score in a group of five identical agents. Then, in order of decreasing performance, we find the Inter + intragroup competition (1) experiment closely followed by the three other experimental conditions. This suggests that the presence of intragroup competition tends

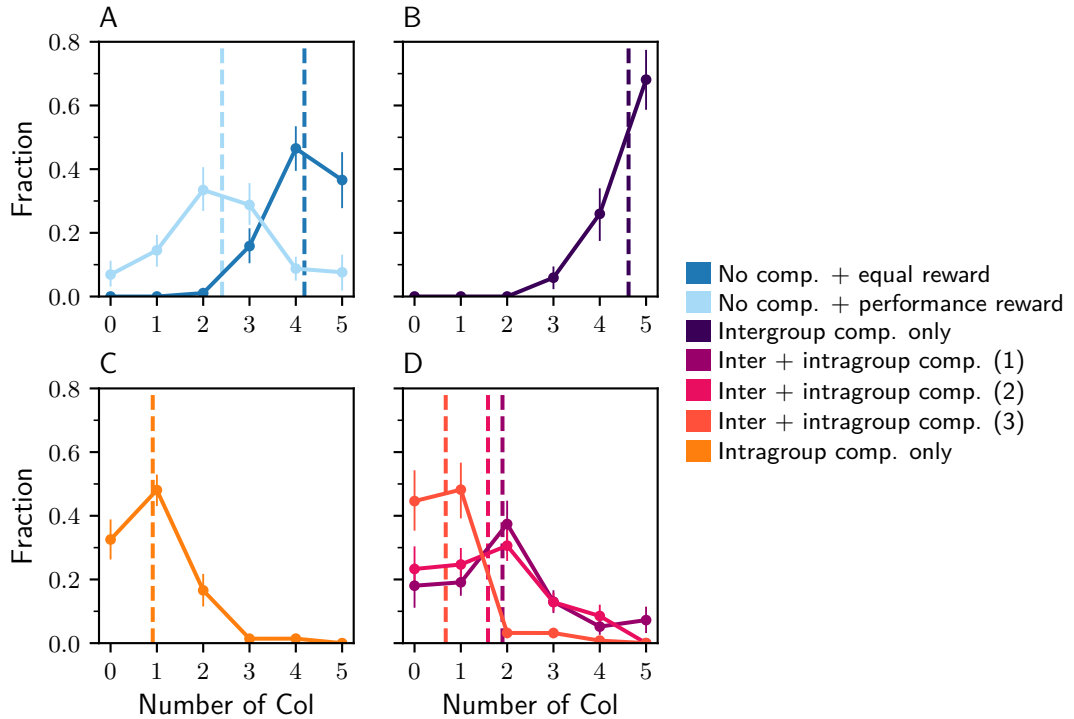


Figure 7.2: **Distribution of collaborators.** Distribution of the number of collaborators in a group for (A) the two conditions without competition, (B) the condition with intergroup competition only, (C) the condition with intragroup competition only, and (D) the three conditions with both intragroup and intergroup competition. The dashed vertical line corresponds to the mean number of collaborators in each experiment.

to lower the performance of the group, due to the competition between the members of the group.

The individuals' rating behavior can be classified into three distinct groups: collaborators, neutrals, and defectors, each reflecting varying degrees of collaboration (see Section 5.4.3). Collaborators assign low ratings to cells with low values and high ratings to those with high values, while defectors exhibit the opposite pattern. Neutrals consistently provide similar ratings irrespective of the cell values. To quantitatively assess the prevalence of these behavioral profiles in each experimental condition, we exploit the methodology of Section 5.4.3 to obtain the fractions of each behavioral profile (see Table 7.1) and the distribution of the number of collaborators in the groups (see Figure 7.2). Notably, the experiment with intergroup competition only stands out with the highest fraction of collaborators (93%), even higher than in the absence of any competition (No comp. + equal reward: 84% of collaborators). This significant level of collaboration between group members is in line with the high scores shown in Figure 7.1A.

Interestingly, in the condition with no explicit competition and a reward proportional to the score, only 48% of the participants adopted a collaborative behavior, while 41% adopted a neutral behavior, and 11% even acted as defectors (see Table 7.1). This finding indicates that the mere notion of a score, which allows participants to compare their

7.3. Performance and Behavioral Profiles of the Participants

		INTERGROUP COMPETITION	
		NO	YES
INTRAGROUP COMPETITION	NO	No comp. + equal reward 84% – 13% – 3%	Intergroup comp. only 93% – 7% – 0%
		No comp. + performance reward 48% – 41% – 11%	
	YES	Intragroup comp. only 18% – 43% – 39%	Inter + intragroup comp. (1) 38% – 41% – 21%
			Inter + intragroup comp. (2) 32% – 42% – 26%
			Inter + intragroup comp. (3) 13% – 46% – 41%

Table 7.1: **Distribution of behavioral profiles.** Percentage of collaborators (green), neutrals (brown), and defectors (red) in the seven experimental conditions, sorted by the presence or absence of intragroup and intergroup competition.

performance to that of others, can introduce an implicit element of competition, even in the absence of explicit competition between participants. As already noted, several participants confirmed this hypothesis by reporting a desire to be the top performer in the group, even in the absence of competition, and although some of them realized that this approach could have negatively impacted their own score.

In the four experimental conditions with intragroup competition, the fraction of collaborators is significantly lower than in the three conditions already addressed, and the fraction of neutrals and defectors is higher. This seems to indicate that intragroup competition favors deception (as seen in Chapter 5). However, the fraction of collaborators and defectors is not the same in all of these conditions. In the Inter + intragroup competition (1) and (2) there are fewer defectors and more collaborators than in the two other conditions. The next section will further develop the origin of these differences.

7.4 Model for Predicting Individual Behavioral Profiles

7.4.1 Three quantities characterizing the payment schemes

We now consider linear regression models, similar to the ones exploited in Section 6.5.2, to predict the average fraction of collaborators, neutrals, and defectors based on the payment schemes used in an experimental condition. We introduce three natural quantities characterizing the payment schemes and which characterize the type and level of competition among the participants. These three quantities will be used as features of the linear regression models:

- **Intra:** The first feature is the standard deviation of the rewards in a group, averaged over both groups. This feature naturally relates to an assessment of the competition within a group (intragroup competition). We will refer to this feature as “Intra”.
- **Inter:** The second feature is the difference between the mean reward in the winning group and in the losing group, where the winning group is the group with the highest group score (average score of its members). This quantity can be interpreted as the pressure of intergroup competition on individuals. It will be referred to as “Inter”.
- **Rank:** The third feature, dubbed “Rank”, is the difference between the mean reward of the five highest-ranked players (ranked by their score) and the mean reward of the five lowest-ranked players. Except for conditions Intergroup comp. only (see Equation (7.1) below) and Intragroup comp. only, the Rank feature coincides with the difference between the mean of the five highest rewards and the mean of the five lowest rewards. This feature relates to the explicit and implicit global competition resulting from the payment schemes and captures some elements of the intragroup *and* intergroup competition.

To compare these three features for different payment schemes with different mean rewards (ranging from 10€ to 15€ in the experiments described here), we normalize them by the mean reward for the corresponding payment scheme.

7.4. Model for Predicting Individual Behavioral Profiles

Although some of these features can be readily computed from the sole knowledge of the payment structure, some others require knowing the actual ranking (according to their score) of the players within a group. However, we aim to design features that can be evaluated before conducting an experiment. Thus, we design a null model of a competition where the ranked players according to their individual scores are randomly distributed in the two groups, hence neglecting correlations between the actual scores and the rank of the players. Within this simple model, the score of a group is simply the opposite of the sum of the ranks of its members. For instance, if the composition of the two groups in terms of ranked players is given by, Group 1 = {2, 3, 4, 6, 10}, and Group 2 = {1, 5, 7, 8, 9}, the two groups have respective scores, $S_1 = -25$, and $S_2 = -30$, and here, Group 1 will be considered the winner. For example, in the Intergroup comp. only condition, the players of Group 1 will receive a 20€ reward, and the players of Group 2 a 10€ reward. Then, the rewards of the players P ranked by their individual score are given by

$$\{P_1 = 10, P_2 = 20, P_3 = 20, P_4 = 20, P_5 = 10, \\ P_6 = 20, P_7 = 10, P_8 = 10, P_9 = 10, P_{10} = 20\}, \quad (7.1)$$

which shows that, in this condition, the player P_1 with the best score/rank does not necessarily obtain the highest reward, and the player P_{10} with the worst score/rank can receive the maximum reward. In this example, one finds Rank = $16 - 14 = 2$ (mean of the top row of Equation (7.1) minus the mean of the bottom row), before normalization by the mean reward in this condition, equal to 15. Then, the three features Intra, Inter, and Rank are computed according to their definition, by averaging the results of 100,000 fictitious games produced by our simple null model. The important point here is that we do not need to measure these features in the actual experiments, and only use the reward scheme to evaluate them with our model. In other words, if a new experimental condition is designed, the features can be evaluated with the sole knowledge of the payment structure, and the linear regression models defined below can be used to predict the fraction of the three behavioral profiles.

However, there is one exception to this possibility to compute the features a priori: the No comp. + performance reward condition, where the real rewards are proportional to the participants' scores, and are not fixed in advance like in all other experimental conditions. In this case, we computed the three features using our simple model, where the players' ranked rewards are the mean ranked rewards obtained by the actual participants in the experiment. These mean values, ranked in decreasing order, are given by, {13.38€, 13.13€, 12.63€, 11.75€, 11.13€, 10.88€, 10.63€, 9.88€, 9.50€, 8.00€}.

Table 7.2 presents the values of the three normalized features for each condition. In the No comp. + equal reward experiment, all features are zero, reflecting the absence of competition and incentives. In the Intergroup comp. only condition, Intra is null and Inter is maximal. Conversely, in the Intragroup comp. only condition, Inter is null and Intra is maximal. In the No comp. + performance reward condition, where there is no explicit competition, the three features are small but nonzero, reflecting the fact that participants could perceive some level of competition, as observed in practice. Moreover, the three Inter + intragroup comp. conditions have low values of Inter and high Intra and Rank

	Inter	Intra	Rank
■ No comp. + equal reward	0	0	0
■ No comp. + performance reward	0.08	0.14	0.24
■ Intergroup comp. only	0.67	0	0.16
■ Inter + intragroup comp. (1)	0.14	0.29	0.61
■ Inter + intragroup comp. (2)	0.15	0.26	0.48
■ Inter + intragroup comp. (3)	0.12	0.26	0.33
■ Intragroup comp. only	0	0.31	0.43

Table 7.2: **Values of the potential features.** Values of the three different normalized features (Inter, Intra, and Rank) in all experimental conditions.

values. Finally, the Intragroup comp. only condition has a zero Inter value, as expected, and the highest Intra value.

7.4.2 Application of the model

Let us now apply the linear regression model, detailed in Section 6.5.2, to the seven experimental conditions studied in this chapter. We have seven different and independent experimental conditions, and two independent variables (C and D , while $N = 1 - C - D$), representing fourteen independent measurements to be explained by the linear regression model. The number of unknown parameters is $2 \times f$, where the number of features f can go from 1 to 3, depending on the number of features used.

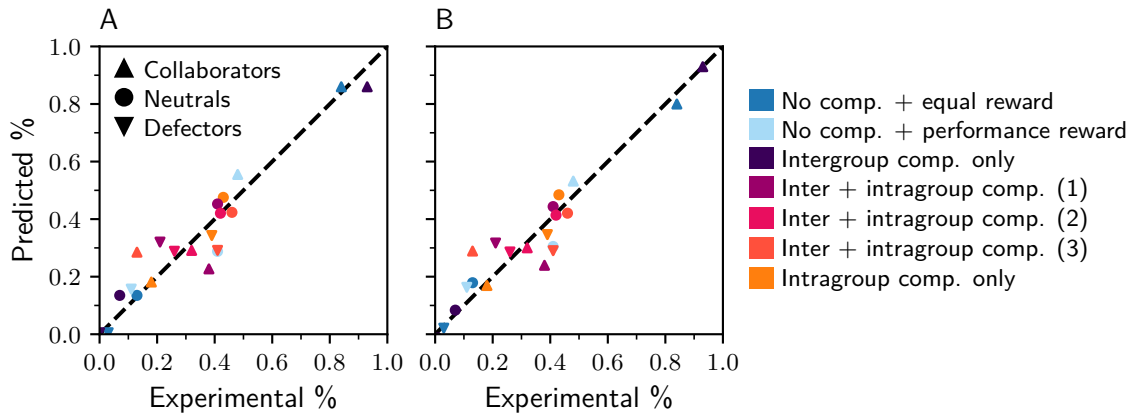


Figure 7.3: **Performance of the Intra model and the Intra + Inter model.** Scatter plot of the predicted fractions of collaborator, neutral, and defectors as a function of the experiment fractions in each experiment for (A) the model with only Intra as feature and (B) the model with Intra and Inter as features. The dotted line represents perfect predictions.

Among the linear models incorporating only one feature among Inter, Intra, and Rank, the model with Intra stands out as a good predictor for the fraction of the three behavioral

7.4. Model for Predicting Individual Behavioral Profiles

profiles, with an error, $E = 0.18$. Figure 7.3A illustrates that although the model performs well, it struggles to accurately discriminate fractions between 20% and 40%. The observed performance can be attributed to Intra's ability to assess both intragroup and intergroup competition. Indeed, when only intergroup competition is present, Intra is equal to 0. However, this feature does not distinguish the Intergroup comp. only condition from the No comp + equal reward condition.

When adding the Inter feature in the Intra model, there is only a very slight improvement, with an error, $E = 0.17$ (compare Figure 7.3B to Figure 7.3A). In fact, the Inter feature is dominated by its high value for the Intergroup comp. only condition (see Table 7.2), and its inclusion only helps to improve the prediction of the Intra model for this condition. Indeed, the model now distinguishes this condition from the No comp + equal reward condition (compare the position of the two triangles in the upper-right of Figures 7.3A and 7.3B). In the Intra + Inter regression model, we also find that the coefficients for the Inter feature are typically ten to twenty times smaller than the ones for Intra. Moreover, in the model with all three features, the coefficients of the Inter feature remain very small, and their sign *changes* compared to the regression model using Intra and Inter. Therefore, the Inter feature does not have any significant explanatory power, in the framework of a linear model.

Let us now address the linear model using the two features Inter and Rank. This four-parameter model, leads to excellent predictions for the fractions of the behavioral profiles, with a reduced error, $E = 0.11$ (see Figure 7.4). The resulting regression reads

$$\begin{cases} C_{\text{pred}} = \mu_C + \sigma_C \left(-1.52 \widehat{\text{Intra}} + 0.64 \widehat{\text{Rank}} \right), \\ N_{\text{pred}} = \mu_N + \sigma_N \left(+1.26 \widehat{\text{Intra}} - 0.38 \widehat{\text{Rank}} \right), \\ D_{\text{pred}} = \mu_D + \sigma_D \left(+1.65 \widehat{\text{Intra}} - 0.84 \widehat{\text{Rank}} \right), \end{cases} \quad (7.2)$$

where $\widehat{\text{Intra}}$ and $\widehat{\text{Rank}}$ are the standardized value of Intra and Rank, and $\mu_C = 0.47$, $\mu_N = 0.33$, $\mu_D = 0.20$, and $\sigma_C = 0.29$, $\sigma_N = 0.15$, $\sigma_D = 0.15$ are the mean fractions of collaborators, neutrals, and defectors, and their standard deviation. Equation (7.2) shows that the Rank feature is two to three times less important than the Intra feature, although including it significantly reduces the error, by improving the predictions of the Intra model for fractions between 20% and 40% (compare Figure 7.3A and Figure 7.4).

Developing Equation (7.2), we obtain

$$\begin{cases} C_{\text{pred}} = 0.47 - 0.43 \widehat{\text{Intra}} + 0.18 \widehat{\text{Rank}}, \\ N_{\text{pred}} = 0.33 + 0.18 \widehat{\text{Intra}} - 0.05 \widehat{\text{Rank}}, \\ D_{\text{pred}} = 0.20 + 0.25 \widehat{\text{Intra}} - 0.13 \widehat{\text{Rank}}. \end{cases} \quad (7.3)$$

In this model, and as expected, an increase in Intra results in a decrease in the fraction of collaborators and an increase in the fractions of neutrals and defectors. Conversely, an increase in Rank increases the fraction of collaborators, while the fractions of neutrals and defectors decrease. In a sense, the Rank feature is better suited to assessing the effective intergroup competition than the Inter feature solely based on the group reward. This Inter




Figure 7.4: **Performance of the Intra + Rank model.** Scatter plot of the predicted fractions of collaborator, neutral, and defectors as a function of the experiment fractions in each experiment. The dotted line represents perfect predictions.

feature is so dominated by the condition with pure intergroup competition that it does not efficiently distinguish the six other conditions.

7.5 Discussion

This chapter presents the results and analyses of seven experimental conditions with different payment schemes, each varying the nature and intensity of competition in the game. These experiments include scenarios with no competition, only intragroup or intergroup competition, and combinations of both.

The analysis of the competitive conditions shows that the payment scheme with intergroup competition favors cooperative behavior within a group, while the one with intragroup competition favors deceptive behavior. This is consistent with the findings of Chapter 5, where I emphasized that competition within a group tends to promote deception. Conditions that combine both types of competition exhibit levels of cooperation between these two extremes. These results are in agreement with the predictions of theoretical multiscale models of the evolution of cooperation, which are not typically based on experimental data. When intergroup competition dominates intragroup competition, it favors cooperative behavior within groups. For instance, in the Inter + intragroup competition (1) condition, while this is not explicitly stated, there is a strong incentive for collaboration within groups since collaborative behavior could lead individuals to be ranked within the first five. This incentive is reflected in the fact that 38% of the participants exhibit a cooperative behavior, compared to only 13% in the Inter + intragroup (3) condition.

Furthermore, in the absence of competition and in the presence of equal rewards, coop-

7.5. Discussion

erative behavior tends to predominate within the group. Interestingly, the introduction of a score coupled with differential rewards, even in a non-competitive environment, leads to a decrease in cooperative behavior. This implies that the quantification of performance and the ability to measure this performance allows for comparison and promotes competition even in the absence of formal competitive structures.

To quantitatively characterize payment schemes, three key features are introduced: *(i)* the standard deviation of rewards within groups, which we relate to an assessment of the intragroup competition; *(ii)* the absolute difference in mean reward between the two groups, which assesses the direct intergroup competition; and *(iii)* the difference between mean rewards of the five highest-ranked and five lowest-ranked players, which quantifies the global competition among the ten participants, and which capture some elements of both intragroup and intergroup competition.

To better understand the impact of payment schemes on individual behavior, we constructed a linear regression model that integrated these three features to predict the proportion of each behavioral profile. We find that *(i)* is the feature with by far the best predictive power, already being able to differentiate the seven conditions according to their intragroup and intergroup competition. Indeed, the linear model based on the sole feature *(i)* already gives good results. Ultimately, the second most relevant feature is *(iii)*, and the model based on *(i)* and *(iii)* leads to an excellent prediction of the fraction of the three behavior profiles. Finally, we find that feature *(ii)*, which accesses the intergroup competition, exhibits only a weak correlation with the data. This feature is dominated by a single condition (the Intergroup comp. only condition) and is unable to distinguish between other conditions. Ultimately, The linear model incorporating features *(i)* and *(iii)* shows that a decrease in *(i)* or an increase in *(iii)* correlates with an increase in the number of collaborators and a decrease in the number of defectors.

Our analysis differs from the traditional perspective of the field, which often categorizes competition solely as a combination of intragroup and intergroup competition. Our study suggests that this Manichean view can fail in practice to capture the subtle interplay between intragroup, intergroup, and a more global notion of competition.

A potential bias in our experimental design lies in the fact that the participants do not have access to the other group's performance throughout the entire session. Consequently, participants have no way of estimating their individual rewards before the end of the session. This design choice may lead some participants to act as if they were competing only against members of their own group, without considering the possibility of cooperating to collectively outperform the other group. This limitation introduces the potential for individual strategies to be influenced by incomplete information about overall group performance, which may affect the observed behaviors and the cooperative dynamics during the games. These results constitute a strong motivation for the LPT-CRCA-TSE collaboration to conduct further experiments based on the STIGMER game, but where, at the end of each game, the players of both groups would be shown their score and rank in their group (as in the conditions studied in this thesis) *and* the performance of the members of the other group.

A missing experiment that would be interesting to do in the future is an experiment

Chapter 7. Impact of Intragroup and Intergroup Competition

without competition, but where participants are paid proportionally to the group score. In this way, the participant's reward depends on the group's performance. This experimental condition would allow us to better quantify the effect of intergroup competition on both performance and cooperative behavior. Indeed, we have seen that in the experiment where there is no competition and participants are paid according to their individual scores, there is some defection, even though there is no need to deceive in the absence of competition. Thus, it is plausible that the observed cooperation in the presence of intergroup competition is not due solely to intergroup competition, but rather to the interdependence of individuals' rewards on those of their group members.

Conclusion

Contents

8.1 Summary of Main Results	137
8.1.1 A game for studying stigmergy	138
8.1.2 Cooperation through stigmergic processes	138
8.1.3 A model of human behaviors	139
8.1.4 Impact of the cooperation level of the group	140
8.1.5 Impact of the type of competition	141
8.2 Limitations	142
8.2.1 Non independence of the games in an experimental session	143
8.2.2 Other means to tune the competition	143
8.2.3 Deception is (too) easy	143
8.2.4 Classification of neutrals	143
8.3 Perspectives	144
8.3.1 Evolution of cooperation	144
8.3.2 Knowledge of the performance of both groups	145
8.3.3 Adaptation in dynamic landscapes	145
8.3.4 No private memory	145
8.3.5 Varying group sizes	146
8.3.6 Analysis of the questionnaires	146

The final chapter of this Ph.D. thesis presents a comprehensive synthesis and discussion of the research findings. Additionally, it highlights the inherent limitations of the study and explores potential future perspectives that deserve consideration.

8.1 Summary of Main Results

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 [Gloag et al., 2013, 2015; Khuong et al., 2016]. This coordination mechanism, known as stigmergy, is widely observed in animal societies and was introduced in Chapter 1. Stigmergy is a process where individuals leave traces in the environment that guide and stimulate the actions of the same or different individuals. In humans, the substantial expansion of the Internet over the past thirty years, combined with applications that

depend on digital traces left willingly or unwillingly by users, has led to the emergence of novel forms of stigmergic processes. An illustration of these processes can be seen in rating-based recommender systems, which recommend content based on the ratings left by previous users.

Although digital traces play a pivotal role in shaping decision-making processes, their impact is not yet fully understood. Therefore, it is important to investigate the influence of digital traces on both individual and collective behaviors [Karanasios et al., 2013; Golder and Macy, 2014; Loh and Kanai, 2016; Firth et al., 2019].

8.1.1 A game for studying stigmergy

To gain a deeper insight into the impact of digital traces on human behavior, Chapter 2 introduces an experimental setup in the form of an online multiplayer game called STIGMER, for which we designed a versatile admin and player interfaces. The game has been used as a platform for conducting experiments involving human participants throughout my Ph.D.

In this game, groups of individuals leave and use digital traces in an information search task, implementing a 5-star rating system. This system is similar to the ones used by many online services such as Amazon, TripAdvisor, or eBay, where users can evaluate products, services, or sellers.

In the STIGMER game, groups of individuals engage in an information search task using a 5-star rating system reminiscent of popular online services such as Amazon, TripAdvisor, or eBay. Participants rate hidden values within a grid, aiming to identify cells with the highest values using only indirect information provided by colored traces generated by collective ratings. This controlled environment facilitates quantitative analysis of individual and collective behavior, and allows manipulation and examination of the combined effects of intragroup and intergroup competition on cooperation.

While we acknowledge that our experimental setup may not capture all the intricacies of real-world rating systems, it does share similarities with them. It both involves exploration of available options (cells in our experiment; products for an online store), which are significantly influenced by their current ratings, and involves participants rating the selected options, which promotes dynamic evolution of ratings.

The flexibility of this experimental setup is remarkable, offering extensive customization options. Parameters such as the number of players, rounds, turns, grid size, values, and shuffling can be controlled, as well as the presence or absence of scoring and the method of calculating scores. The STIGMER game also includes the possibility of having an arbitrary number of preprogrammed bots to play against the human participants. This versatility allows the game to be used in a wide variety of scenarios and to study very different experimental conditions.

8.1.2 Cooperation through stigmergic processes

The first part of this thesis focuses on the study of cooperation and deception through stigmergic processes. In particular, we investigate the conditions under which cooperation

8.1. Summary of Main Results

emerges through digital traces and indirect interactions.

Chapter 5 presents, analyzes, and compares two experiments, without and with competition among participants. The results indicate that groups of individuals can effectively use colored traces derived from their ratings to coordinate their search and collectively discover cells with the highest values in a grid of hidden numbers. These traces serve as a form of long-term collective memory, documenting past actions taken by the group [Thierry et al., 1996; Baltzersen, 2022]. Combined with individual short-term memory of cell values visited, these traces influence participants' choices of cells to ultimately explore.

Furthermore, upon closer examination of individual behaviors, variations in the utilization of digital information become apparent. Three distinct categories of behavior emerge based on their degree of cooperation: collaborators, who rate proportionally to cell value; defectors, who rate inversely to values; and neutrals, who provide the same average rating regardless of cell value.

The comparison of the competitive and non-competitive experiments reveals significant differences in the usage and reliability of the trace. In the absence of competition, individuals tend to exhibit cooperative behavior, resulting in the generation of reliable digital traces. However, the introduction of competition drastically reduces cooperation, leading to more neutral and defector-like behaviors. Consequently, some individuals choose to withhold private information and may even spread misinformation, leading to a decrease in the quality of the social information (the colors on the grid).

8.1.3 A model of human behaviors

In addition to the data analysis, our main focus was the design of a model that captures the observed human behaviors within the experimental setup.

Chapter 4 introduces a stochastic agent-based model that simulates individual behaviors in the STIGMER game. The model features agents governed by strategies, each comprising two independent components: the visit strategy, which determines cell selection, and the rating strategy, which dictates cell rating. The model's flexibility is reflected in its simple parameterization, which allows for the creation of agents with different behaviors.

The model's strength lies in its relative simplicity, which offers a clear interpretation and understanding of the underlying mechanisms that govern the visit and rating behaviors in the STIGMER game. Despite its simplicity, the model is grounded in real-world observations of human behavior, making it a relevant and effective tool for studying stigmergic processes.

The model serves multiple purposes. Firstly, it enhances our understanding and offers a concise and precise representation of individual and collective human behaviors in the STIGMER game. Additionally, simulations using agents with human-like behaviors serve as predictive tools for human behavior in unexplored situations. Furthermore, the model is exploited in controlling the behavior of social bots, facilitating experiments in a fully controlled environment alongside human participants.

Although the model's simplicity has advantages, it also leads to some limitations. In particular, the model omits a possible explicit time dependence in the parameters governing visit and rating strategies. The model simplifies the time-dependence aspect by relying

on the inherent time-dependence of cell colors and their three best-discovered cells. Although the model does not consider time-dependent visit parameters for simplicity and dataset constraints, it is noteworthy that even with its imperfections, the model's agreement with experimental results remains within acceptable limits. This level of agreement suggests that the model captures the essential dynamics of actual visit and rating behaviors, despite some limitations in reproducing certain observables.

8.1.4 Impact of the cooperation level of the group

Although the earlier discussed model effectively captures human behavior within the game, it fails to explain how individuals adapt their exploration and rating strategies based on the behaviors of other players in the group. Specifically, the model does not predict the composition of the group in terms of collaborators, neutrals, and defectors, depending on the observed fractions in experiments, which are fed by hand in the model.

Chapter 6 addresses this gap by delving into the possible cues influencing a human participant's adoption of a specific behavioral profile. The chapter introduces experiments in which a human interacts with four social bots controlled by the aforementioned model. By manipulating the behavior of these bots, the setup provides a controlled environment for studying human behavior and its adaptation under various conditions. The utilization of collaborator and defector bots allows for a precise control of the level of cooperation within a game, facilitating a comprehensive analysis of human behavior and social interactions in diverse collaborative contexts.

The results indicate that human participants consistently outperform the simple bots by adopting predominantly deceptive or neutral behaviors. This outcome was expected due to the predefined strategies of the bots, which restrict their ability to adapt to human behavior.

Furthermore, a correlation has been found between the distribution of behavioral profiles and the level of cooperation of the bots. As the fraction of collaborator bots increases, we observe a higher probability for human players to act as defectors. Indeed, some participants exploit the collaborative bots to manipulate traces and lower the performance of the bots (who cannot retaliate by changing their strategy). These results indicate that, in a competitive context, some individuals may not reciprocate reliable social information, even when the majority of their group are cooperators.

To understand the basis for individual choices of behavioral profiles, we designed a linear regression model that uses cues available to players during the game. These cues include:

- (i) an assessment of whether colored cells correspond to high- or low-value cells implicitly measuring the degree of collaboration in the group;
- (ii) an assessment of the number of different significantly rated cells, which provide information about the legibility and discriminative nature of the social information;
- (iii) an assessment of a player's performance/rank relative to other players indicating the effectiveness of their strategy.

8.1. Summary of Main Results

According to this statistical model, an increase in the average values of highly colored cells or an increase in the number of different rated cells is correlated with an increase in collaborators and a decrease in defectors. Additionally, neutral or defector behaviors are generally associated with good ranks, while bad ranks are correlated with an increased likelihood of collaborative behavior, suggesting that some participants switch their strategy to improve their ranking.

Chapter 6 presents a significant finding of the study, demonstrating the influence of social bots on human behavior in scenarios involving intragroup competition. By combining the model of human behavior that controls the behavior of bots with the regression model that predicts the behavioral profile of humans, it is possible to design social bots that adapt to the collaborative nature of their environment. Alternatively, such bots could promote various behaviors in humans, and in particular, the deliberate manipulation of human behavior by social bots presents new opportunities to foster cooperative behavior in groups.

8.1.5 Impact of the type of competition

Our analysis has primarily focused on experiments conducted without competition or with competition within the same group. However, it is important to acknowledge that real-world scenarios often involve multiple groups competing within the same environment, creating a complex hierarchy of multilevel competition [Wilson, 2015; Cooney et al., 2023]. This is particularly significant as prevailing theories on the evolution of cooperation propose that the presence of multiscale competition is imperative for cooperation to evolve. In essence, intragroup competition tends to favor deception, which theoretically should prevent cooperation in a purely selfish world. However, in nature, various groups of individuals often compete collectively. When intergroup competition surpasses intragroup competition, cooperation can emerge within groups.

Chapter 7 introduces seven experimental conditions with distinct payment schemes, each varying the nature and intensity of competition among ten participants in our experiments (two groups of five players). The experiments encompass scenarios with no competition, solely intragroup or intergroup competition, and combinations of both. The objective is to examine the impact of intragroup and intergroup competition on individual performance and behaviors. This chapter presents experimental results that provide empirical insights into the often theoretical domain of multiscale theories of evolution.

The analysis of the competitive conditions indicates that a payment scheme with intergroup competition encourages cooperative behaviors within a group, while pure intragroup competition promotes deceptive behaviors. Conditions that combine both types of competition exhibit levels of cooperation between these two extremes. These findings are consistent with the predictions of theoretical multiscale models of the evolution of cooperation.

To quantitatively characterize the competition in the payment schemes, three key features are introduced:

- (i) the standard deviation of rewards within groups, which we relate to an assessment

of the intragroup competition;

- (*ii*) the absolute difference in mean reward between the two groups, which assesses the direct intergroup competition;
- (*iii*) the difference between mean rewards of the five highest-ranked and five lowest-ranked players, which quantifies the global competition among the ten participants, and which captures some elements of both intragroup and intergroup competition.

To better understand the impact of payment schemes on individual behavior, we constructed a linear regression model that integrated these three features to predict the proportion of each behavioral profile. We find that (*i*) which assesses the intragroup competition is the feature with by far the best predictive power, already being able to differentiate the seven conditions according to their intragroup and intergroup competition. Indeed, the linear model based on the sole feature (*i*) already gives good results. Ultimately, the second most relevant feature is (*iii*), and the model based on (*i*) and (*iii*) leads to an excellent prediction of the fraction of the three behavior profiles. Finally, we find that feature (*ii*), which access the pure intergroup competition, exhibits only a weak correlation with the data. This feature is dominated by a single condition (the Intergroup comp. only condition), and is unable to distinguish between other conditions. Ultimately, the linear model incorporating features (*i*) and (*iii*) shows that a decrease in (*i*) or an increase in (*iii*) correlates with an increase in the number of collaborators and a decrease in the number of defectors.

The data analysis supports the basis of theoretical multiscale models that aim at explaining the evolution of cooperation. However, our regression model reveals a divergence from the prevailing perspective in the field, which often only emphasizes intragroup and intergroup competition. The current perspective is inadequate in the practical context of our experiments, as it does not consider the complex interplay between intragroup, intergroup, and a more global competition, beyond just the notion of groups. Additionally, our research indicates that individual behavior is not solely impacted by the nature of the competition. The capacity to measure and compare performance between individuals can also affect their behavior. This performance-based comparison can induce competition even in the absence of formal competitive structures. Hence, real-life competition scenarios often involve complexities beyond the sole notions of pure intragroup and intergroup competition.

8.2 Limitations

The experimental and theoretical methodology and tools exploited in this thesis establish the foundations for understanding stigmergic interactions in digital environments, shedding light on the relationships between competition, cooperation, deception, and decision-making. However, the present section will also address certain limitations associated with our study.

8.2. Limitations

8.2.1 Non independence of the games in an experimental session

In our experiments, participants played ten games in groups of five. This method allowed us to study changes in behavior within a given environment. However, the ten games played by a group are not independent due to significant correlations in behavior between games performed by the same group. This limits the number of independent experiments that we have. An alternative would be to shuffle the participants between the two groups for each game of the experimental session. This would increase the independence of individual games, when desired.

8.2.2 Other means to tune the competition

In Chapters 5 to 7, our analysis is based on varying type and level of competition controlled by different payment schemes. In particular, the conditions studied in Chapter 7 exploit a manipulation of the payment scheme to tune the intragroup and intergroup competition. However, the participants did not have an explicit knowledge of the performance of the members of the other group, and we clearly missed a way to better infuse the notion of intergroup competition among the players. This limitation constitutes a strong motivation for the LPT-CRCA-TSE collaboration to conduct further experiments based on the STIGMER game, but where, at the end of each game, the players of both groups would be shown their score and rank in their group (as in the conditions studied in this thesis) *and* the performance of the members of the other group.

8.2.3 Deception is (too) easy

In our setup, white cells indicate either unvisited cells or those that have been visited but rated with zero stars. As a result, an individual who encounters a high-value cell can discreetly assign it a zero-star rating, allowing the cell to go unnoticed. This principle is often exploited by defectors. This contrasts with typical online rating systems, where the distinction between no ratings and low ratings is clear. Criticism could be raised regarding the excessive simplicity of this strategy, and its nonconformity with actual rating systems.

It should be emphasized that our experimental setup is not a direct representation of a real-life rating system. Rather, it is a tool that allows for the study of cooperation and defection in human groups, inspired in part by such systems. However, it would be possible to address this issue and make deception more difficult while also moving closer to real-life rating systems. One possible approach is to differentiate between never-visited cells and those that have been visited but rated zero by using different colors. For instance, non-visited cells could be colored in gray. Another option is to set the minimum rating to one instead of zero, which is a common practice in online services.

8.2.4 Classification of neutrals

In our experiments, we classified individuals into three distinct behavioral profiles based on their mean ratings: collaborators, who exhibit cooperative behavior; defectors, who engage in deceptive strategies; and neutrals, a category that proves to be more diverse

and challenging to characterize uniformly. It is noteworthy that the proportion of neutral individuals remains consistently high across most experiments, typically around 40%.

Within the neutral category, subgroups with varying rating strategies were identified. Some neutrals consistently assign the same rating, mostly zero, one or five. Others demonstrate random rating behaviors, introducing an element of unpredictability.

Currently, our classification method groups together all of these neutral profiles, without the ability to distinguish between them. Further investigation and refinement of the classification method is necessary to uncover the subtleties within the neutral category. Additionally, adding a time-dependent dimension to the behavioral profiles could provide a more comprehensive understanding of how individuals' strategies develop throughout the game. Note that this task could benefit from a detailed analysis of the questionnaire filled by the participants at the end of a one-hour session (see Appendix A.2). During the thesis, we did not really exploit this survey, except to check that the participants consistently declared that they understood the rules and principles of the game and enjoyed it (except when playing against 4 defector bots!).

8.3 Perspectives

The STIGMER game introduced and analyzed in this thesis is a versatile tool for exploring cooperation and deception through stigmergic interactions within human groups. Its high degree of customization allows for easy implementation of new variations, each providing unique insights. In addition, the game can be made publicly available on the Internet, allowing for large-scale experiments to be conducted.

This section outlines various potential extensions of the work presented in this thesis are outlined. While some are conceptual ideas, others have already undergone some experimentation and/or analysis.

8.3.1 Evolution of cooperation

Multilevel selection is believed to have played a crucial role in shaping cooperative behavior within human groups [Traulsen and Nowak, 2006; Luo, 2014; Simon et al., 2013; Cooney, 2019; Wilson, 2015].

The hypothesis suggests that when intergroup competition exceeds intragroup competition in intensity, intragroup cooperation becomes more likely. Chapter 7 discusses experiments exploring different types and levels of competition to analyze their impact on cooperation. However, it is also possible to simulate the evolution of cooperation and deception in our experimental framework.

The agent-based model presented in Chapter 4 which is used to reproduce human-like behavior in the STIGMER game can be used to simulate games with agents of varying levels of cooperation. An evolutionary algorithm could be used to evolve groups of agents within specific competitive contexts. This approach could provide valuable insights into how cooperation evolves over time in dynamic, competitive settings.

Although I began this line of research during my third year of Ph.D. study, time constraints led to its interruption. However, this is something I would like to investigate

8.3. Perspectives

further.

Moreover, it could also be possible to perform real-life experiments with individuals performing repetition of the game, possibly mixing groups, like in repeated prisoner dilemma experiments that are conducted in experimental economics.

8.3.2 Knowledge of the performance of both groups

During our competitive experiments, the participants only received information about the performance of the members of their own group, without any information about the other group. This incomplete information could have limited their perception of the intergroup nature of the experiment, and would have certainly limited their ability to react to the other group displaying a better performance, for instance, by starting to collaborate more with the members of their group. To better understand this phenomenon, additional experiments should be conducted where participants are shown the performance of both groups at the end of a game. This can be achieved by displaying either the ranking of all ten players, or both the within-group ranking and the between-group ranking.

8.3.3 Adaptation in dynamic landscapes

All experiments conducted during my Ph.D. research were performed on games with a static grid of values. This setup, similar to the assumption of constant quality in real-world rating systems, oversimplifies the dynamic nature of many scenarios. Indeed, the quality of objects, products, or services can vary due to factors such as defects, varying product quality, or changing restaurant experiences.

To address this limitation, introducing noise to the values in our experimental setup could better emulate real-world fluctuations. Each visit to a cell would yield a different value, requiring players to make multiple visits for an accurate estimation. This dynamic element might encourage increased cooperation among group members to collectively refine their estimates of cell values.

The evolution of the “quality” of the cells over time can be simulated by varying the values of the cells over time. This approach may discourage revisits to previously identified high-value cells and encourage continuous exploration of the grid, potentially impacting player behavior. This setup would enable the study of how human groups adapt to changing landscapes and their resistance to noisy information.

In these situations, it would be crucial to include a temporal aspect to the ratings. Ratings should have a duration, causing older ratings to gradually lose influence over time. This fading of old ratings would reduce their impact on cell color, with recent ratings carrying more weight. Note that we have already tested rating evaporation on a static grid but did not observe any significant performance changes.

8.3.4 No private memory

In the STIGMER game, participants have access to both private/individual and shared/collective memory. Individual memory includes knowledge about the location and values of

previously visited cells, while collective memory involves shared information represented by the colors of the cells.

During my Ph.D., we conducted a series of experiments in which we eliminated the private memory, forcing individuals to rely solely on shared memory. To achieve this, we informed the participants that the grid cells and their attached color would be shuffled between each round. As a result, participants could only rely on the ratings to identify the localization of high-value cells. The results of these experiments have not yet been fully analyzed, but it is already clear that collaboration is significantly fostered by the elimination of private information (65% of collaborators, compared to only 18% in the competitive experiment addressed in Chapter 5).

8.3.5 Varying group sizes

Chapter 7 introduces an experimental condition without competition, where participants' rewards are based on their individual performance (score). Although these experiments were conducted in groups of five, additional experiments were conducted in groups of two and ten, which have not yet been analyzed. Note that in all experimental conditions, the first two games were played alone and were analyzed in Section 5.6.

This setting was designed to promote collaboration among participants to achieve the highest possible score. However, our findings, in groups of five, show that the presence of the score and the ability to compare with others affect participants, resulting in less collaborative behavior. Thus, the experimental design offers an opportunity to examine the effect of group size on both individual and collective performance in a group.

Furthermore, this experimental design is a valuable tool for exploring how individuals influence each other in groups of varying sizes. Indeed, in groups of two, individuals can easily identify each other's cooperative and defection tendencies, and we can expect to observe tit-for-tat behaviors.

8.3.6 Analysis of the questionnaires

In several conducted experiments, participants were required to fill out a questionnaire at the end of a one-hour session. The questionnaire (see Appendix A.2) asked participants about their understanding of the STIGMER game, their strategies for visiting and rating, their trust in social information, their qualitative strategy, and their overall enjoyment of the game. Although the questionnaire contains valuable information, it has not yet been analyzed in details, and only confirmed the good understanding of the STIGMER game rules and principles, and the fact that the participants reported enjoying it. The analysis of this survey could lead to a better understanding of the participants' behavior and help determine if their self-reported behavioral profile matches their actual actions.

Appendices

Documents Related to the Experiments

Contents

A.1 Information Note and Informed Consent Form	150
A.2 Questionnaire	152
A.3 JSON File	153
A.4 CSV File	155
A.5 Ethics Approvals	156

A.1 Information Note and Informed Consent Form

Notice d'information et consentement éclairé

Titre du projet

Analyse expérimentale et modélisation des choix collectifs dans les groupes humains

Chercheurs titulaires responsables scientifiques du projet

- Dr. Clément Sire, Laboratoire de Physique Théorique, Centre National de la Recherche Scientifique, Université de Toulouse Paul Sabatier, Toulouse, (clement.sire@univ-tlse3.fr)
- Dr. Guy Theraulaz, Centre de Recherches sur la Cognition Animale, Centre de Biologie Intégrative, Centre National de la Recherche Scientifique, Université de Toulouse Paul Sabatier, Toulouse (guy.theraulaz@univ-tlse3.fr)

Autres chercheurs impliqués dans ce projet

- Thomas Bassanetti, Laboratoire de Physique Théorique & Centre de Recherches sur la Cognition Animale, Centre de Biologie Intégrative, Centre National de la Recherche Scientifique, Université de Toulouse Paul Sabatier, Toulouse
- Dr. Ramón Escobedo, Centre de Recherches sur la Cognition Animale, Centre de Biologie Intégrative, Centre National de la Recherche Scientifique, Université de Toulouse Paul Sabatier, Toulouse
- Dr. Adrien Blanchet, Toulouse School of Economics, INRA, Université de Toulouse Capitole, Toulouse
- Stéphane Cezera, Toulouse School of Economics, INRA, Université de Toulouse Capitole, Toulouse

Lieu de recherche

Behavioral Economics Lab, Toulouse School of Economics, Université de Toulouse Capitole, Toulouse

But du projet de recherche

Analyse des choix collectifs dans des groupes humains.

Ce que l'on attend de vous (méthodologie)

Vous allez participer à un jeu qui consiste à découvrir les valeurs maximales dans un tableau de nombres.

A.1. Information Note and Informed Consent Form

Vos droits de vous retirer de la recherche en tout temps

1. Vous participez à cette recherche sur la base de votre volontariat.
2. Vous pouvez interrompre votre participation à tout moment.
3. Sans avoir à en justifier les raisons et sans aucun préjudice.

Vos droits à la confidentialité et au respect de la vie privée

1. Les données obtenues seront traitées de manière totalement anonyme.
2. Votre identité sera masquée par un numéro de participant aléatoire.
3. Aucun autre renseignement ne sera dévoilé qui puisse révéler votre identité.
4. Toutes les données seront gardées dans un endroit sécurisé et seuls les responsables scientifiques y auront accès.

Risques possibles

Cette recherche n'implique aucun risque ou inconfort.

Diffusion

Les résultats de cette recherche sont susceptibles d'être diffusés dans des colloques et d'être publiés dans des revues académiques sans qu'il ne soit possible d'identifier votre performance personnelle.

Vos droits de poser des questions en tout temps

Vous pouvez poser des questions à propos de la recherche en tout temps en communiquant avec les investigateurs principaux (Clément Sire : clement.sire@univ-tlse3.fr ou Guy Theraulaz : guy.theraulaz@univ-tlse3.fr).

Consentement à la participation

En signant le formulaire de consentement, vous certifiez que vous avez lu et compris les renseignements ci-dessus, qu'on a répondu à vos questions de façon satisfaisante et qu'on vous a avisé que vous étiez libre d'annuler votre consentement ou de vous retirer de cette recherche en tout temps, sans préjudice.

À remplir par le participant

J'ai lu et compris les renseignements ci-dessus et j'accepte de plein gré de participer à cette recherche.

Nom, Prénom :

Email :

Date :

Signature :

Signature de l'investigateur :

A.2 Questionnaire

Questionnaire anonyme de fin de session

Numéro de session :

Numéro de joueur :

Pour chacune des affirmations suivantes, indiquez sur une échelle de 1 à 5 à quel point vous êtes d'accord avec elle (1 : pas du tout d'accord ; 2 : pas d'accord ; 3 : ni d'accord ni pas d'accord ; 4 : d'accord ; 5 : tout à fait d'accord).

1. Vous pensez avoir bien compris les règles du jeu et votre objectif en tant que joueur.

1	2	3	4	5
---	---	---	---	---

2. Quand vous avez dû choisir une case à ouvrir, vous avez choisi cette case :

a. En utilisant sa couleur

1	2	3	4	5
---	---	---	---	---

b. Parce que vous connaissiez déjà sa valeur

1	2	3	4	5
---	---	---	---	---

c. Parce que vous ne l'aviez pas encore visitée

1	2	3	4	5
---	---	---	---	---

3. Quand vous avez dû noter une case ouverte, vous avez utilisé :

a. Sa valeur

1	2	3	4	5
---	---	---	---	---

b. Sa couleur

1	2	3	4	5
---	---	---	---	---

c. Si vous n'avez utilisé ni sa valeur ni sa couleur, expliquez en quelques mots ce qui a déterminé votre notation :

--

4. Les notes que vous avez attribuées aux cases que vous avez visitées vous ont servi à :

a. Vous souvenir des cases de grandes valeurs

1	2	3	4	5
---	---	---	---	---

b. Éviter de revisiter des cases de faibles valeurs

1	2	3	4	5
---	---	---	---	---

c. Aider les autres joueurs

1	2	3	4	5
---	---	---	---	---

d. Tromper les autres joueurs

1	2	3	4	5
---	---	---	---	---

5. Notez votre confiance en les marques qui ont été laissées collectivement par les autres membres de votre groupe.

1	2	3	4	5
---	---	---	---	---

6. Avez-vous apprécié le jeu ?

1	2	3	4	5
---	---	---	---	---

A.3 JSON File

```
1 {
2   "id": 1674568506077,
3   "session": {
4     "phaseId": "2023_01_TSE",
5     "sessionId": 10,
6     "groupId": "A",
7     "session_type": "R2_intra",
8     "date": "2023-1-24"
9   },
10  "players": [
11    {
12      "id": "P01",
13      "age": 22,
14      "gender": "male"
15    },
16    {
17      "id": "P02",
18      "age": 21,
19      "gender": "female"
20    },
21    {
22      "id": "P03",
23      "age": 20,
24      "gender": "female"
25    },
26    {
27      "id": "P04",
28      "age": 19,
29      "gender": "male"
30    },
31    {
32      "id": "P05",
33      "age": 22,
34      "gender": "female"
35    }
36  ],
37  "game": {
38    "timeStart": "14:55:37",
39    "timeEnd": "14:58:12",
40    "game_type": "R2_intra",
41    "altGameName": "Group_5",
42    "numberRounds": 20,
43    "numberCellsOpenedPerRound": 3,
44    "numberPlayers": 5,
45    "botsList": [],
46    "evaporation": "1000",
47    "randomMS1": "on",
```

Appendix A. Documents Related to the Experiments

```
48     "randomMS2": null,
49     "bots": [],
50     "rule": {
51         "_id": "63c52ecdbb7ee4436c9d938b",
52         "rule": 2,
53         "ruleName": "rule_2",
54         "maxNumberOfStarsPerRound": 15,
55         "coefRemainingStars": 0,
56         "coefValueTimesStars": null,
57         "coefValue": "on"
58     },
59     "versus": false,
60     "map": {
61         "_id": "63c52ecdbb7ee4436c9d938c",
62         "mapName": "random_1",
63         "mapType": "random",
64         "map": [
65             [44, 3, 4, 1, 2,20,20, 1, 2, 1, 1, 3, 2, 0, 3],
66             [ 1, 1, 2, 1, 4, 5,19, 8, 1, 2, 4,85, 1, 8,45],
67             [ 2, 2,99, 1,71, 3, 2, 1,20, 2, 1, 2, 3, 2,72],
68             [84,86, 0, 3, 6, 3, 0, 2, 7, 2, 7, 3, 2, 2, 2],
69             [ 4, 1, 2, 1, 4, 1, 2, 2,86, 3, 2,53, 8, 3, 3],
70             [ 2, 1,44, 3, 1, 1, 2, 0, 2, 6,28, 1, 1, 7, 7],
71             [27, 3,43, 0, 6, 2, 1, 2, 2,12, 0,22, 2,53, 1],
72             [ 0, 3,51, 2, 4, 2, 1, 3,11, 1, 0, 3, 0, 1, 2],
73             [ 1, 2,20, 1,13, 1,71, 3, 3, 2, 0, 1,14, 3, 1],
74             [ 3, 9, 2, 3,27, 1, 2, 8, 0, 5, 6, 2, 1, 6,20],
75             [ 2, 2, 1, 6, 3,21, 2, 0, 1, 1, 0, 3, 3, 3, 1],
76             [ 2,27, 4,43, 1,19, 0, 2, 3, 0, 0, 6, 0, 1,72],
77             [ 0, 1, 2,14, 2, 4, 3, 3,24,53, 2, 4, 2, 8, 0],
78             [ 2,45, 1, 0,13, 2, 4, 4, 6,44, 3,13, 3, 0, 0],
79             [ 1, 1, 2, 9, 3,12,24, 1, 7, 0, 0,46,24, 3, 1]
80         ]
81     }
82 }
83 }
```

Listing A.1: **JSON file example.** Example JSON file from an experiment. This file contains detailed data about the session, the players, and the game. This file contains all the data needed to replicate the experiment. The corresponding CSV file is shown in Table A.1.

A.4. CSV File

A.4 CSV File

round	playerId	mapX	mapY	value	numberStars	score
1	P03	7	2	1	0	1
1	P04	14	0	3	0	3
1	P02	4	8	13	0	13
1	P03	13	10	3	0	3
1	P04	0	0	44	0	44
1	P01	3	11	43	1	43
1	P05	1	8	2	0	2
1	P03	2	7	51	1	51
1	P02	11	11	6	3	6
1	P04	0	14	1	0	1
1	P01	7	7	3	5	3
1	P05	12	6	2	5	2
1	P02	11	4	53	1	53
1	P01	11	12	4	5	4
1	P05	7	3	2	5	2
2	P03	2	7	51	0	51
2	P04	0	0	44	0	44
2	P02	11	4	53	1	53
2	P01	3	11	43	0	43
2	P03	3	2	1	0	1
2	P04	2	7	51	0	51
2	P05	11	11	6	5	6
2	P02	11	12	4	4	4
2	P03	13	2	2	0	2
2	P05	11	12	4	5	4
2	P04	3	11	43	0	43
2	P02	5	12	4	0	4
2	P01	4	1	4	1	4
2	P05	3	11	43	0	43
2	P01	14	0	3	1	3

Table A.1: **CSV file example.** Table corresponding to the first thirty lines of the CSV file of the game corresponding to the JSON file above (Listing A.1). The CSV file contains all of the players' actions during the game, i.e., their visits and their ratings.

A.5 Ethics Approvals



Toulouse, 24 September 2021

Research ethics approval: “Experimental analysis and modeling of collective choices in human groups”

To whom it may concern:

The TSE Research Ethics Committee for Experimental Research has evaluated the request submitted by Clément Sire and Guy Theraulaz for ethics approval of their study “Experimental analysis and modeling of collective choices in human groups”, with Adrien Blanchet, Stéphane Cezera, Ramón Escobedo and Thomas Bassanetti.

The Committee has carefully studied the documents provided in the application and approved the project today.

Yours sincerely,

A handwritten signature in black ink, appearing to read 'Tiziana Assenza'.

Tiziana Assenza
(Chair, TSE Research Ethics Committee for Experimental Research)

A.5. Ethics Approvals



A l'attention de
THERAULAZ Guy

À Toulouse, le 27 octobre 2021

Affaire suivie par :
Christina WATKINS
CER-DRDV
Courriel : bureau-cer@univ-toulouse.fr
Tél. : 05 61 10 80 30

Objet : Avis du comité pour le projet resousmis 2021-425

Titre du projet soumis : Analyse expérimentale et modélisation des choix collectifs dans les groupes humains

Porteur de projet : THERAULAZ Guy, laboratoire CRCA - CBI, UT3

Monsieur,

Compte tenu des éléments fournis dans votre demande, le Comité d'Ethique pour les Recherches de l'Université de Toulouse émet l'avis suivant : **Avis Favorable avec recommandations.**

Recommandations : préciser le lieu de stockage des formulaires de consentement éclairé (conseil du CER: placer l'enveloppe dans une armoire sécurisée).

Nous rappelons, par ailleurs, qu'il relève de la responsabilité des chercheurs de se conformer à leurs obligations légales notamment en ce qui concerne les aspects d'homologation du lieu de recherche ou RGPD : Règlement Général sur la Protection des Données.

Le Comité d'Ethique rappelle au porteur de projet qu'il doit tenir compte des conditions sanitaires actuelles et mettre en œuvre un protocole sanitaire adapté en conformité avec les recommandations des tutelles.

Nous restons à votre disposition pour toute question.

Les membres du bureau CER.

Pr Jacques Py



Dr Rémi Capa



Pr Maria Teresa Munoz Sastre



Résumé en français

Contents

B.1 Introduction	159
B.2 Dispositif Expérimental	161
B.3 Impact de la Présence ou Absence de Compétition sur la Coopération Stigmergique	162
B.3.1 Expériences	163
B.3.2 Comportements collectifs	164
B.3.3 Comportements individuels	164
B.3.4 Modèle de comportement des joueurs	167
B.3.4.1 Définition du modèle	167
B.3.4.2 Prédiction du modèle	167
B.4 Impact de la Composition du Groupe sur les Comportements Individuels et la Coopération Stigmergique	168
B.4.1 Expériences	168
B.4.2 Impact des <i>bots</i> sur la performance des participants	169
B.4.3 Modèle des stratégies de visites et de notations	169
B.4.4 Prédiction des profils comportementaux des participants humains	171
B.5 Impact de la Compétition Intragroupe et Intergroupe sur la Coopération Stigmergique	172
B.5.1 Expériences	172
B.5.2 Performance et profils comportementaux des participants	174
B.5.3 Modèle de prédiction des profils comportementaux individuels	174
B.6 Conclusion	177

B.1 Introduction

L'intelligence collective désigne la capacité des communautés à relever ensemble des défis, à prendre des décisions et à générer des connaissances qui dépassent les capacités individuelles [Bonabeau, 2009]. Elle repose sur deux phénomènes : l'information sociale et l'influence sociale. L'information sociale est la connaissance et les indices que les individus d'un groupe partagent entre eux [Camazine et al., 2001]. L'influence sociale est le processus par lequel les individus ajustent leurs attitudes, leurs comportements ou leurs opinions en réponse à des situations sociales [Cialdini and Goldstein, 2004].

Le concept d'intelligence collective trouve une application pratique dans des mécanismes tels que la stigmergie [Grassé, 1959], qui se retrouve couramment dans les sociétés animales et humaines [Theraulaz and Bonabeau, 1999]. La stigmergie est un mécanisme de coordination dans lequel les traces laissées par les individus dans un milieu guident et stimulent leurs actions ultérieures. Parmi les différents exemples de stigmergie, on peut citer le comportement de recherche de nourriture des fourmis [Goss et al., 1989], la construction de nid par les guêpes [Theraulaz and Bonabeau, 1995] et la création de lignes de désir (sentier tracé graduellement par érosion à la suite du passage répété de piétons, cyclistes ou animaux) [Helbing et al., 1997]. Cependant, la digitalisation de la société humaine a donné naissance à une nouvelle forme de stigmergie qui s'appuie sur des traces numériques [Van Dyke Parunak, 2006]. Parmi les exemples de ces processus stigmergiques figurent notamment les systèmes de notation et de recommandation qui utilisent largement les traces numériques laissées par leurs utilisateurs [Lü et al., 2012].

La coopération est le processus par lequel des groupes travaillent ensemble pour obtenir des avantages communs ou mutuels [Axelrod, 2006]. La prévalence de la coopération parmi les espèces sociales soulève des questions sur son évolution dans un monde dans lequel la sélection naturelle peut favoriser l'égoïsme [Wilson, 2015]. Plusieurs mécanismes ont été proposés pour expliquer cette évolution, notamment la sélection de parentèle, la réciprocité directe, la réciprocité indirecte et la sélection de groupe [Nowak, 2006]. Cependant, des mécanismes stigmergiques peuvent aussi permettre de mieux comprendre la coopération par le biais d'interactions indirectes [Chiong and Kirley, 2012]. La compétition entre individus, qui s'exerce à différents niveaux, est un autre facteur déterminant dans le développement de la coopération [Toma and Butera, 2015]. En effet, la rivalité entre les individus et les groupes peut être un catalyseur d'amélioration, mais elle peut aussi inciter à des comportements contraires à l'éthique ou trompeurs [Gilbert, 2006].

De nombreuses stratégies ont été proposées et testées pour améliorer l'intelligence collective dans les groupes humains. Par exemple, il est possible d'y parvenir en changeant la taille [de Condorcet, 1785] ou la composition du groupe [Page, 2008] ou même la structure du réseau d'interactions [Newman, 2003]. Une autre option consiste à améliorer l'échange d'informations sociales en augmentant le contrôle sur celles-ci et en réduisant le bruit [Kahneman et al., 2021]. Les "nudges" [Thaler and Sunstein, 2008] et les "bots" [Jayles et al., 2020a] sociaux peuvent aussi être utilisés pour influencer les comportements.

Dans la société numérique actuelle, les informations sociales sous forme de traces numériques ont un impact significatif sur la prise de décision humaine [Chavalarias, 2022]. Il est donc crucial de mieux comprendre l'impact sociétal de ces traces, en particulier si l'on considère leur vulnérabilité à la manipulation. Dans ce contexte, cette thèse de doctorat vise à étudier comment les humains utilisent et interagissent avec les traces numériques, en mettant l'accent sur les conditions de coopération et l'influence de la compétition sur leur utilisation.

B.2. Dispositif Expérimental

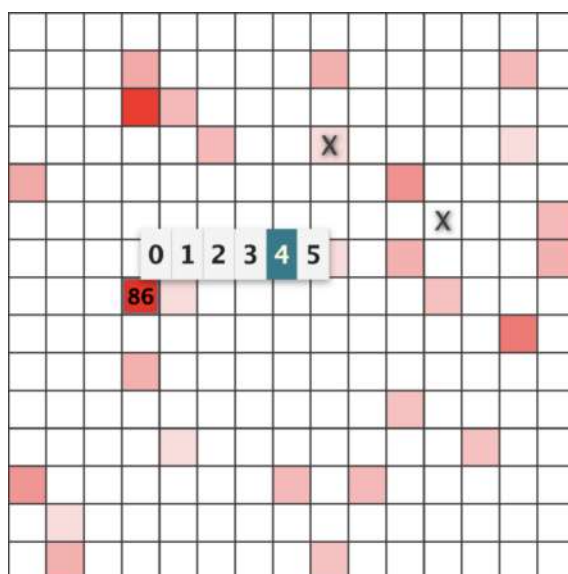


FIGURE B.1 : **Exemple de tableau.** Capture d'écran du tableau tel que le voit un joueur au tour 8. On peut y voir le tableau avec les cellules colorées, les croix noires qui marquent les cellules déjà visitées pendant le tour, ainsi que la valeur de la cellule actuellement sélectionnée, accompagnée de l'interface de sélection de la notation.

B.2 Dispositif Expérimental

Le jeu STIGMER est un jeu multijoueur dans lequel les participants s'engagent dans une recherche collective d'informations dans un environnement numérique. Conçu à l'origine par Guy Theraulaz et Clément Sire, plusieurs extensions du jeu STIGMER ont été développées et étudiées expérimentalement et théoriquement au cours de ma thèse de doctorat. Le jeu STIGMER reproduit certains aspects des systèmes de notation par étoiles utilisés sur de nombreuses plateformes sur Internet telles qu'Amazon, TripAdvisor et eBay, où les utilisateurs notent les produits, les services ou les vendeurs afin de trouver le meilleur choix possible. L'objectif principal pour les joueurs est l'identification des meilleures options parmi celles disponibles, avec l'aide uniquement d'interactions indirectes sous forme de notations collectives.

Dans ce jeu, les joueurs doivent identifier les cellules avec des valeurs élevées dans un tableau de 225 cellules (15 lignes et 15 colonnes) chacune contenant une valeur cachée (voir Figure B.1). Les valeurs des cellules sont des entiers compris entre 0 et 99, avec une prédominance de petites valeurs, et sont aléatoirement distribués dans chacune des cellules du tableau. Les cellules représentent les options disponibles, et leur valeur correspondent à leur qualité intrinsèque.

Chaque jeu comprend vingt tours successifs. Au cours de chaque tour, chaque joueur doit visiter et évaluer successivement trois cellules distinctes. Lorsqu'un joueur découvre la valeur cachée d'une cellule, il doit la noter de zéro à cinq étoiles. Le tour se termine lorsque tous les membres du groupe ont visité et évalué trois cellules différentes.

Au début du tour suivant, la couleur de chaque cellule du tableau est mise à jour en

fonction de la fraction d'étoiles utilisées pour noter la cellule depuis le début de l'expérience, c'est-à-dire le nombre d'étoiles dans la cellule divisé par le nombre total d'étoiles dans toutes les cellules. L'échelle des couleurs varie entre le blanc (0%) et le noir (100%) en passant par un gradient de nuances de rouge. Ainsi, les cellules qui ont reçu la plus grande fraction d'étoiles depuis le début de l'expérience seront clairement visibles par tous les individus appartenant au même groupe. Ces couleurs agissent comme une mémoire collective à long terme pour le groupe, qui est mise à jour à chaque tour.

Le jeu peut être joué avec ou sans score. Lorsqu'il y a un score, celui-ci augmente à chaque tour des valeurs des cellules visitées durant le tour, indépendamment de leur notation. Les participants voient leur score en temps réel tout au long du jeu, ce qui leur permet de suivre leur progression. De plus, à la fin du jeu, un tableau affiche les scores et les classements finaux de tous les joueurs du jeu, ce qui leur permet d'évaluer leur performance par rapport à celle des autres.

Afin de faciliter l'expérimentation, le jeu a été implémenté sous la forme d'une application web qui est accessible depuis n'importe quel appareil et n'importe quelle localisation. Cette application en ligne offre la flexibilité d'exécuter plusieurs variations de jeu, chacune avec des différences de règles telles que le nombre de joueurs, le nombre de tours, le calcul des scores, etc. Elle fournit aux participants une interface conviviale pour une expérience de jeu agréable et nous donne un accès pratique pour configurer et suivre la progression des expériences.

Pendant mes études doctorales, nous avons mené plusieurs séries d'expériences (voir Figure B.2) qui étaient divisées en sessions. Chaque session durait une heure, durant laquelle les participants (regroupés en un ou deux groupes de cinq joueurs) jouaient à une douzaine de répétitions du jeu STIGMER. À la fin de la session, les participants recevaient une récompense monétaire pouvant être corrélée ou non à leur performance dans la recherche des cellules de plus forte valeurs. En changeant cette récompense, nous pouvons modifier les incitations des joueurs et créer différent niveau de compétition (intragroupe et/ou intergroupe) entre les joueurs. Par conséquent, cette configuration expérimentale est un excellent outil pour étudier les processus de coopération et de compétition au sein et entre les groupes.

B.3 Impact de la Présence ou Absence de Compétition sur la Coopération Stigmergique

Cette section a deux objectifs. Premièrement, elle vise à présenter et à comprendre comment les individus utilisent les interactions indirectes (stigmergiques) pour interagir dans le contexte de notre expérience de recherche d'informations. Deuxièmement, cette section explore l'impact d'un contexte compétitif ou non compétitif sur la coopération et la tromperie entre les individus et la manière dont ils échangent et utilisent les informations sociales, intégrées dans les traces de leurs actions passées, pour réaliser la tâche de recherche d'informations.

B.3. Impact de la Présence ou Absence de Compétition sur la Coopération Stigmergique



FIGURE B.2 : **Photographies des expériences.** Photographies des expériences menées au Laboratoire d'économie comportementale de Toulouse School of Economics. (A) Salle d'attente. (B) Salle dans laquelle les règles sont expliquées aux participants. (C) Salle d'expérimentation. (D et E) Participants jouant au jeu STIGMER. (F) Paiement d'un participant.

B.3.1 Expériences

Cette section présente et compare les résultats de deux expériences basées sur le jeu STIGMER utilisant des schémas de récompense différents pour explorer l'impact de la compétition sur la coopération, ainsi que les stratégies individuelles de visite et de notation.

Pour les deux expériences, chaque session expérimentale regroupe dix participants. Chaque participant commence par jouer deux parties de vingt tours, seul, dans l'objectif de se familiariser avec l'interface web et d'intégrer les règles. Ensuite, les participants sont répartis de manière aléatoire en deux groupes indépendants de cinq joueurs et jouent dix parties de vingt tours dans ces mêmes groupes.

Dans les deux expériences, le but de jeu est le même : trouver les cellules de valeur élevée. Cependant, il y a une différence dans le score et la rémunération finale, ce qui conduit une expérience à avoir une compétition entre les participants et l'autre à ne pas avoir de compétition.

Dans l'expérience non-compétitive (ci-après appelée Règle 1), il n'y a pas de score et tous les participants sont payés le même montant de 10 €. Dans ce contexte, les actions des participants (visites et notations) n'ont aucun impact sur leur paiement, assurant ainsi qu'il n'y a pas de compétition entre les joueurs.

Dans l'expérience compétitive (ci-après appelée Règle 2), le score des participants correspond à la somme des valeurs des cellules qu'ils ouvrent. À la fin d'une session expérimentale, les scores obtenus par un participant dans l'ensemble des jeux sont combinés

pour obtenir son score total. Ensuite, dans chaque groupe, les participants sont classés en fonction de leur score total respectif et sont payés selon leur classement. Le joueur arrivé en première position reçoit 20 €, le deuxième reçoit 15 €, et les trois autres joueurs reçoivent chacun 10 €. Par conséquent, les participants sont incités à obtenir le meilleur score, créant ainsi une compétition au sein des groupes.

B.3.2 Comportements collectifs

Afin d’analyser les performances des individus et des groupes, ainsi que la dynamique de l’exploration collective et des notations dans les deux règles, nous introduisons un ensemble d’observables. Pour résumer, nous avons introduit : le score des individus (S) ; le score des groupes (\hat{S}) ; le score moyen sur un jeu ($\langle S \rangle$) ; la valeur moyenne des cellules pondérée par la fraction d’étoiles ou la fraction de visites au tour t ($p(t)$ et $q(t)$) ou jusqu’au tour t ($P(t)$ et $Q(t)$) ; le nombre effectif de cellules (“Inverse participation ratio” ; IPR) sur lesquelles les étoiles et les visites sont distribuées au tour t ($\text{IPR}p(t)$ et $\text{IPR}q(t)$) et jusqu’au tour t ($\text{IPR}P(t)$ et $\text{IPR}Q(t)$) ; la fidélité F , qui quantifie si la distribution des étoiles ou des visites dans chaque cellule coïncide avec la distribution réelle des valeurs des cellules ($F(\mathbf{V}, \mathbf{P}(t))$ et $F(\mathbf{V}, \mathbf{Q}(t))$).

L’analyse de ces observables montre que la compétition (Règle 2) offre une incitation plus forte à visiter les cellules de haute valeur par rapport à la Règle 1. Cependant, les individus visitent significativement plus de cellules différentes dans la Règle 1 que dans la Règle 2. En effet, dans la Règle 2, les individus revisitent beaucoup plus de cellules de haute valeur au lieu d’explorer de nouvelles cellules afin de maximiser leur score. Concernant la notation des cellules, les étoiles sont en moyenne sur des cellules de valeurs plus élevées en Règle 1 par rapport à la Règle 2. Cela est dû aux différences de stratégies des individus pour noter les cellules dans les deux règles. En effet, nous verrons plus loin que dans la règle compétitive, certains individus choisissent de donner une notation moyenne ou même faible aux cellules de haute valeur pour éviter de les signaler aux autres membres de leur groupe. Enfin, la fidélité des notations est significativement plus élevée dans la Règle 1, où les participants explorent davantage le tableau et où leurs notations reflètent mieux la valeur des cellules découvertes.

B.3.3 Comportements individuels

Cette section est dédiée à la caractérisation des comportements des individus et leurs stratégies de visite et de notation des cellules, c’est-à-dire la manière dont ils utilisent l’information sociale sous la forme de traces colorées résultant de leurs actions collectives passées.

Dans la Règle 1, où il n’y a pas de compétition, les individus trouvent plus souvent les cellules de haute valeur que ce que prévoit le hasard (voir Figure B.3). Cela montre un effet coopératif entre les participants. En revanche, dans la Règle 2, les individus revisitent fréquemment les cellules à haute valeur pour améliorer leur score, ce qui limite leur capacité à découvrir de meilleures options. Ainsi, bien qu’en moyenne les participants ouvrent des cellules de valeurs plus élevées dans la Règle 2, cela se fait au détriment

B.3. Impact de la Présence ou Absence de Compétition sur la Coopération Stigmergique




FIGURE B.3 : **Probabilité de trouver les cellules ayant les valeurs les plus élevées.** (A) Probabilité de trouver la meilleure cellule, de valeur 99. (B) Probabilité moyenne de trouver une des quatre cellules dont les valeurs sont 86 (), 85, ou 84. (C) Probabilité moyenne de trouver une des quatre cellules dont les valeurs sont 72 () ou 71 (). La Règle 1 non-compétitive est en bleu et la Règle 2 compétitive est en orange. Les points représentent les données expérimentales et les lignes pleines les prévisions du modèle. Les lignes noires en pointillés et en tirets correspondent aux probabilités attendues de deux stratégies de visite différentes : cellules choisies au hasard, et cellules choisies séquentiellement, respectivement.

de l'exploration de nouvelles options. C'est donc en Règle 1 que les cellules de haute valeur sont trouvées. Globalement, ces résultats soulignent l'impact de la compétition sur la stratégie d'exploration des individus, avec moins d'exploration et plus de revisite des options sûres dans un environnement compétitif.

Dans les deux règles, on observe trois motifs de notation émergents parmi les participants lorsqu'ils notent des cellules en fonction de leur valeur. Certains individus notent les cellules de manière proportionnelle à leur valeur (collaborateurs), aidant ainsi les autres membres du groupe à identifier les meilleures cellules. D'autres notent les cellules indépendamment de leur valeur (neutres), ne fournissent pas d'information distinctive aux autres membres du groupe. Finalement, d'autres attribuent des notes inversement proportionnelles à la valeur des cellules (défecteurs ou trompeurs), dissimulant les meilleures cellules et induisant les autres en erreur. Voir Figure B.4 pour la notation moyenne donnée par chaque profil.

Pour quantifier et classifier ces trois profils comportementaux, une fonction affine est ajustée aux notes moyennes de chaque individu en fonction de la valeur de la cellule. En absence de compétition, une large majorité de participants adoptent un profil collaborateur, cependant en présence de compétition le pourcentage de collaborateur est très faible au profit des neutres et défecteurs (voir Table B.1). À l'inverse, les défecteurs ont accès à plus d'information que les autres membres du groupe. Cette asymétrie d'information leur permet de maximiser leurs scores dans un environnement compétitif. En revanche, en absence de compétition, l'intérêt est moindre et on observe moins de défecteurs.

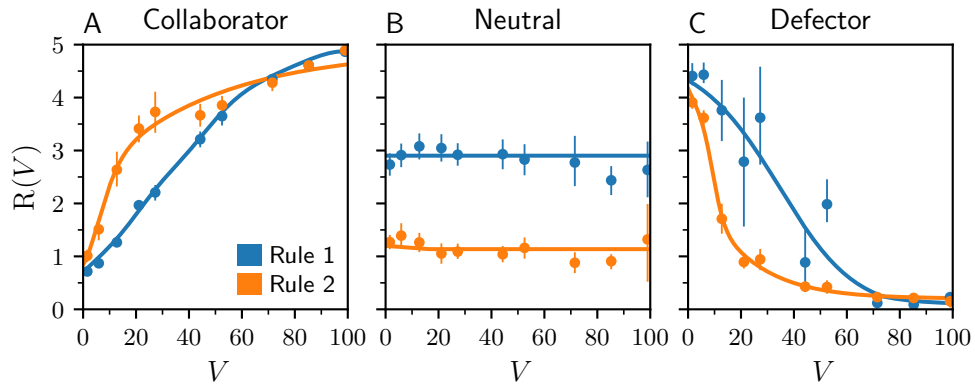


FIGURE B.4 : **Note moyenne donnée par chaque profil comportemental.** Nombre moyen d'étoiles utilisées pour noter les cellules en fonction de la valeur V de la cellule pour (A) les collaborateurs, (B) les neutres et (C) les défecteurs. La Règle 1 est en bleu et la Règle 2 est en orange. Les points représentent les données expérimentales et les lignes pleines les prédictions du modèle.

	Col	Neu	Def
■ Règle 1	84	13	3
■ Règle 2	18	43	39

TABLE B.1 : **Profil comportementaux des participants.** Pourcentage d'individus avec des comportements collaboratifs, neutres et défecteurs dans chaque règle.

B.3. Impact de la Présence ou Absence de Compétition sur la Coopération Stigmergique

B.3.4 Modèle de comportement des joueurs

B.3.4.1 Définition du modèle

Nous introduisons un modèle d'agent stochastique utilisé pour simuler le comportement des joueurs dans le jeu de STIGMER. Le modèle se compose d'agents dotés de différentes stratégies qui régissent leurs actions au cours du jeu. En particulier, la stratégie de chaque agent est décomposée en deux composantes indépendantes : la stratégie de visite, qui régit la sélection des cellules, et la stratégie de notation, qui régit la notation des cellules.

Au cours d'une partie, les joueurs ont accès à des informations collectives et privées. L'information collective est représentée par la couleur des cellules du tableau, c'est-à-dire les notes, tandis que l'information privée correspond à la connaissance que les joueurs ont des cellules qu'ils ont visitées précédemment. La stratégie de visite prend en compte les deux types d'informations dont disposent les agents. Cela leur permet soit de revisiter des cellules déjà visitées lors du tour précédent, en utilisant des informations privées, soit de visiter d'autres cellules en fonction de leur couleur, en utilisant des informations collectives. Lorsqu'ils utilisent l'information collective, les agents peuvent choisir de se concentrer sur les cellules sombres, en suivant l'information sociale, ou sur les cellules blanches, pour l'exploration des cellules non évaluées. La stratégie de notation attribue une probabilité de donner une certaine note à une cellule visitée, qui dépend généralement de la valeur de la cellule.

Le modèle est contrôlé par des paramètres qui permettent de générer des agents aux comportements différents en changeant la valeur de ses paramètres. Parmi ces variations, on trouve des agents qui explorent beaucoup, d'autres qui évitent l'exploration et d'autres encore qui s'engagent dans des pratiques de notation collaboratives ou trompeuses.

Le modèle est ensuite utilisé pour créer différents types d'agents. Des "agents Mimic" qui imitent le comportement humain, des "agents optimisés" qui sont adaptés à des tâches et à des environnements spécifiques, et des agents dotés de stratégies spécifiques qui sont utilisés pour comprendre l'impact de différentes stratégies sur les performances individuelles et collectives.

Tous les paramètres régissant ces agents sont soit spécifiés manuellement, soit déterminés en fonction des caractéristiques du comportement humain, soit déterminés à l'aide de méthodes d'optimisation (en particulier, Monte-Carlo).

B.3.4.2 Prédiction du modèle

En calibrant le modèle, nous pouvons reproduire quantitativement la dynamique de l'exploration et des notations collectives, ainsi que les performances individuelles et collectives observées dans les deux expériences. Nous observons alors un bon accord entre l'expérience et le modèle qui reproduit une série d'observables subtils. L'analyse des comportements individuels combinée aux simulations du modèle informatique montre que la compétition renforce le poids de l'information privée (c'est-à-dire la mémoire individuelle des cellules déjà visitées) par rapport à l'information sociale (c'est-à-dire la mémoire collective du groupe représentée sur le tableau colorée partagée) dans le choix des cellules qui sont visitées.

Une analyse plus poussée du modèle indique qu'un effet coopératif induit par la trace émerge dès qu'il existe un niveau minimal de marquage des cellules et que la fidélité des notations augmente avec la coopération. Le modèle montre également que la trace induit une faible coopération, même dans les groupes de défecteurs, à condition qu'ils notent les cellules avec un nombre suffisant d'étoiles, simplement parce qu'ils revisitent les cellules dont les valeurs sont les plus élevées. Dans ce cas, la mémoire individuelle joue un rôle majeur dans la performance collective de ces défecteurs. Par ailleurs, le modèle prédit que l'effet coopératif induit par les traces et la performance moyenne des individus augmente avec la taille du groupe. Cette propriété résulte des interactions stigmergiques entre les individus qui permettent d'amplifier au niveau du groupe l'information sur la localisation des cellules dont les valeurs sont les plus élevées.

Nous avons aussi utilisé le modèle pour trouver des agents optimisés dans différentes situations. Pour ce faire, nous avons utilisé une méthode de Monte-Carlo afin d'obtenir les paramètres du modèle qui caractérisent les stratégies de visite et de notation correspondantes. L'analyse de ces agents optimisés indique que le score maximal est obtenu pour les agents collaboratifs, ce qui suggère que la collaboration intragroupe devrait émerger de la compétition intergroupe. De plus, le modèle prédit également qu'un comportement défecteur émerge lorsqu'un agent vise à optimiser son rang, dans les mêmes conditions que celles de notre expérience.

B.4 Impact de la Composition du Groupe sur les Comportements Individuels et la Coopération Stigmergique

Dans la section précédente, nous avons mis en évidence la coopération à travers des interactions indirectes du jeu STIGMER, en distinguant trois profils comportementaux selon leurs stratégies de notation : les collaborateurs, les neutres et les défecteurs. Bien qu'un modèle ait été développé sur la base de ces profils, il ne permet pas de prédire l'adaptation des comportements des individus en réponse aux autres. Cette section a pour but de comprendre les signaux qui influencent les comportements des participants grâce à des expériences où la composition des groupes est contrôlée.

B.4.1 Expériences

Cette section présente et compare neuf nouvelles conditions expérimentales. Les expériences consistent en une répétition du jeu STIGMER et sont jouées avec un score défini comme la somme des valeurs visitées au cours du jeu. De la même manière que lors de l'étude de la Règle 2, à la fin de la session expérimentale, les cinq participants sont classés en fonction de leur score cumulé sur l'ensemble des jeux et payés en conséquence. Le joueur classé premier reçoit 20 €, le deuxième 15 € et les trois joueurs restants (classés entre 3 et 5) reçoivent 10 € chacun. Les participants sont ainsi incités à obtenir le meilleur score de leur groupe, ce qui crée une compétition entre eux.

Contrairement aux expériences décrites précédemment, où les cinq participants d'un groupe jouent ensemble, dans les expériences présentées ici, chaque participant joue avec

B.4. Impact de la Composition du Groupe sur les Comportements Individuels et la Coopération Stigmergique

quatre *bots* contrôlés par le modèle (voir Section B.3.4.1). Ce dispositif expérimental nous permet d'examiner comment le comportement des *bots* affecte le comportement des participants humains dans un environnement entièrement contrôlé. Pour garantir l'équité du paiement des participants, qui sont classés ensemble mais ne jouent pas dans le même groupe, tous les participants d'une même session expérimentale jouent sur des tableaux identiques, avec le même mélange de valeurs, et contre des *bots* aux stratégies identiques.

Chacune des neuf conditions expérimentales comporte une composition de groupe différente. Les cinq premières conditions testent différentes combinaisons de *bots* collaborateurs et défecteurs, allant d'un scénario avec quatre collaborateurs et zéro défecteur (4 Col - 0 Def) à un scénario avec zéro collaborateur et quatre défecteurs (0 Col - 4 Def). Les trois conditions expérimentales suivantes impliquent que les joueurs interagissent avec quatre *bots* neutres qui vont toujours donner la même note quelle que soit la valeur de la cellule. Trois types de *bots* différents sont testés : des *bots* donnant toujours une note d'une étoile (Const-1), de trois étoiles (Const-3) et de cinq étoiles (Const-5). Dans la dernière condition, les joueurs interagissent avec quatre *bots* qui ont été optimisés pour maximiser le score du groupe (Opt).

B.4.2 Impact des *bots* sur la performance des participants

Quel que soit le niveau de collaboration du groupe, les participants humains ont réussi à surpasser les *bots*. Lorsque le niveau de coopération des *bots* augmente, le score des participants humains augmente également, et ils sont plus susceptibles de trouver les cellules de plus grande valeur sur le tableau. Ces résultats mettent en évidence l'influence de la composition du groupe sur la dynamique de coopération et les performances individuelles. Les participants humains ont obtenu cette meilleure performance en adaptant leur stratégie de visite et de notation au comportement des quatre autres membres de leur groupe (*bots*), ce qui a une incidence sur les informations sociales partagées (les cellules colorées du tableau). Cette possibilité de s'adapter à l'environnement rencontré n'était pas disponible pour les *bots*, dont la stratégie de visite et de notation était fixe pendant toutes les parties, et ils ont été pénalisés par cette asymétrie.

De plus, nous avons découvert une corrélation entre la distribution des profils comportementaux et le niveau de coopération dans le jeu (voir Figure B.5). Plus précisément, nous avons observé que plus le nombre de *bots* collaborateurs augmentait, plus les joueurs humains étaient susceptibles d'adopter un comportement de type défecteur. Il semble que certains participants humains exploitent la nature collaborative du jeu pour manipuler la trace et diminuer les performances des autres.

B.4.3 Modèle des stratégies de visites et de notations

Nous avons ensuite utilisé le modèle décrit précédemment pour modéliser les stratégies de visites et de notations des humains dans ces expériences. En examinant les valeurs des paramètres, nous constatons que le seuil à partir duquel les individus commencent à revisiter les cellules du tour précédent est plus bas lorsque les participants jouent avec de nombreux défecteurs. Ainsi, lorsque les participants jouent avec des défecteurs, ils se

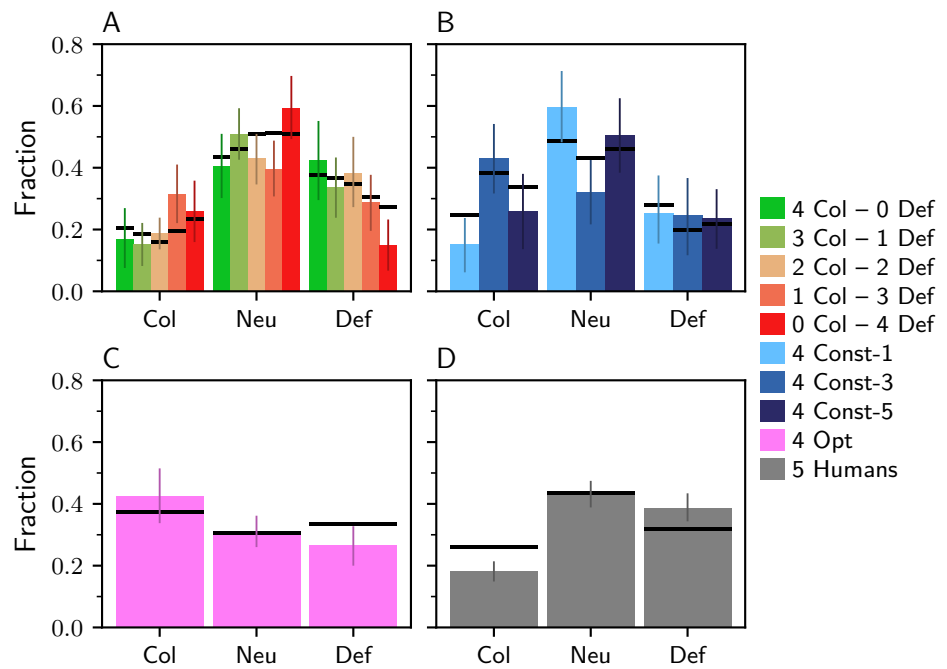


FIGURE B.5 : **Profil comportemental des participants humains.** Fraction de collaborateurs, de neutres et de défecteurs pour les expériences dans lesquelles un participant humain joue avec (A) des *bots* collaborateurs et défecteurs, (B) des *bots* constants et (C) des *bots* optimisés, et (D) pour l'expérience avec cinq participants humains. Les lignes noires correspondent aux prédictions du modèle avec (i), (ii) et (iii).

B.4. Impact de la Composition du Groupe sur les Comportements Individuels et la Coopération Stigmergique

contentent de valeurs plus basses au lieu de continuer à chercher des valeurs plus élevées. En outre, nous constatons que dans les conditions où l'information sociale est fiable (c'est-à-dire que les cellules foncées correspondent à des valeurs plus élevées que les cellules claires), les participants humains ont systématiquement tendance à accorder plus de crédit à la couleur des cellules du tableau. Cela indique que les participants humains sont bien conscients du degré de collaboration des quatre autres membres de leur groupe et qu'ils peuvent adapter leur stratégie de visite et de notation sur la base de cette observation qualitative.

B.4.4 Prédiction des profils comportementaux des participants humains

L'analyse a révélé que les participants fondaient leur choix de profil comportemental sur trois signaux principaux disponibles pendant le jeu :

- (i) une évaluation de la coopération au sein du groupe (valeur moyenne d'une cellule colorée) ;
- (ii) une évaluation de la clarté de la l'information sociale (nombre de cellules différentes qui ont été évaluées) ;
- (iii) une évaluation de l'efficacité de leur stratégie (performance, i.e., rang, par rapport aux autres joueurs du jeu).

Un jeu typique caractérisé par un haut niveau de coopération présente des caractéristiques distinctes : les cellules qui sont évaluées par les joueurs sont principalement celles qui ont des valeurs élevées, et le nombre de cellules colorées reste relativement faible. À l'inverse, dans les situations caractérisées par un manque de coopération, un grand nombre de cellules sont évaluées, mais ces cellules ont tendance à avoir des valeurs faibles ou intermédiaires.

Pour mieux comprendre comment les individus utilisent ces signaux pour prendre des décisions, nous avons construit un modèle linéaire, incorporant les trois signaux décrits ci-dessus pour prédire la proportion de chaque profil comportemental (voir Figure B.6). Ce modèle simple a donné des résultats prometteurs et a offert une interprétation significative du comportement humain observé dans les différentes conditions contrôlées. En particulier, nous avons constaté que parmi les trois signaux, les valeurs des cellules colorées et le nombre de cellules colorées sont plus influents que le classement sur le comportement des participants.

Le modèle linéaire indique qu'une augmentation des valeurs moyennes des cellules très colorées ou une augmentation du nombre de cellules différentes évaluées est corrélée à une augmentation du nombre de collaborateurs et à une diminution du nombre de défecteurs. En outre, le modèle montre que les bons classements sont plus favorables à un comportement neutre ou défecteur, tandis que les mauvais classements sont corrélés à une probabilité plus élevée de comportement collaboratif. Cela suggère que les participants peuvent stratégiquement passer d'un comportement de collaboration à un comportement neutre ou défecteur afin d'améliorer leur classement.

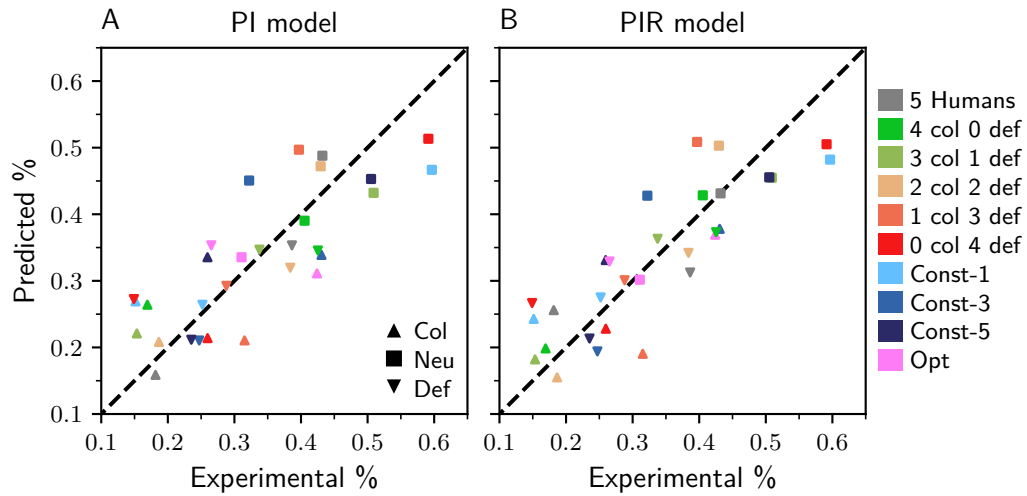


FIGURE B.6 : Performance du modèle avec (i) et (ii) et du modèle avec (i), (ii) et (iii). Fractions prédites de collaborateurs, de neutres et de défecteurs en fonction des fractions expérimentales dans chaque expérience pour (A) le modèle avec (i) et (ii) et (B) le modèle avec (i), (ii) et (iii). La ligne en pointillé représente les prédictions parfaites.

B.5 Impact de la Compétition Intragroupe et Intergroupe sur la Coopération Stigmergique

Jusqu'à présent, notre attention s'est portée sur des expériences se déroulant soit en absence de compétition, soit en présence de compétition au sein d'un même groupe. Cependant, dans la nature, les scénarios sont souvent caractérisés par la présence de multiples groupes en compétition dans un même environnement, générant ainsi une compétition complexe à plusieurs niveaux.

Les modèles multi-échelles de l'évolution de la coopération prédisent que la coopération au sein d'un groupe peut être favorisée lorsque la compétition entre les groupes est plus intense que celle à l'intérieur du groupe [Wilson, 2015; Cooney, 2019]. Pour tester cette hypothèse, nous présentons une série d'expériences examinant les effets de différentes structures de récompense sur la compétition intragroupe et intergroupe, ainsi que leur incidence sur les performances collectives et les comportements individuels.

B.5.1 Expériences

Dans cette section, nous étudions l'impact de sept structures de paiement combinant différents niveaux de compétition intragroupe et intergroupe. Chacune de ces structures de compétition utilise différentes méthodes de paiement pour rémunérer les participants à la fin d'une session expérimentale. En manipulant la structure de paiement, il est possible de créer différents types et niveaux de compétition, tant entre les participants du même groupe (compétition intragroupe) qu'entre différents groupes de participants (compétition intergroupe). Même dans le même type de compétition, la structure de paiement peut être ajustée pour fournir des incitations légèrement différentes aux joueurs.

B.5. Impact de la Compétition Intragroupe et Intergroupe sur la Coopération Stigmergique

Chaque session expérimentale implique dix participants. Tout d'abord, chaque participant joue seul deux parties de vingt tours. Les participants sont ensuite répartis au hasard en deux groupes de cinq joueurs chacun et jouent environ dix parties de vingt tours dans leurs groupes respectifs. À la fin de chaque session expérimentale, les joueurs sont payés selon une structure de paiement spécifique, qui dépend de la condition expérimentale.

Les deux premières structures de paiement ne créent aucune compétition explicite entre les participants, mais correspondent à différentes incitations :

- **Pas de compétition + récompense égale** : Dans cette condition, il n'y a ni score ni classement entre les participants. Chaque participant reçoit une récompense fixe de 10 €. Ici, les individus sont libres de faire ce qu'ils veulent, car leur comportement n'affecte pas leur récompense.
- **Pas de compétition + récompense basée sur la performance** : Dans cette condition, il y a un score, mais les participants ne sont pas classés, et sont payés proportionnellement à leur score individuel. En moyenne la récompense est de 11 €. Ce schéma de paiement encourage les individus à atteindre le score le plus élevé possible, ce qui devrait conduire à de la coopération entre les participants.

Le deuxième type de structure de paiement repose uniquement sur la compétition intragroupe. Il y a compétition entre les membres d'un groupe, mais pas de compétition entre les deux groupes. Les cinq joueurs de chaque groupe sont classés indépendamment en fonction de leurs scores et payés en conséquence :

- **Compétition intragroupe seulement** : Dans chaque groupe, le participant classé premier reçoit 20 €, le participant classé deuxième reçoit 15 €, et chacun des trois autres participants (classés trois à cinq) reçoit 10 €.

Le troisième type de structure de paiement implique uniquement la compétition intergroupe. Il n'y a pas de compétition explicite ou implicite entre les cinq membres d'un même groupe, mais il y a une compétition explicite entre les deux groupes. Les deux groupes sont classés en fonction du score moyen de leurs cinq membres :

- **Compétition intergroupe seulement** : Chaque participant du groupe avec le score le plus élevé (score moyen de ses membres) reçoit 20 €, tandis que chaque membre de l'autre groupe reçoit 10 € chacun.

Le quatrième type de structure de paiement combine à la fois la compétition intragroupe et intergroupe. À la fin de la session expérimentale, les dix joueurs (cinq dans chaque groupe) sont classés ensemble en fonction de leurs scores individuels et récompensés en conséquence. Bien qu'il y ait à la fois une compétition intragroupe et intergroupe dans les trois expériences décrites ci-dessous, les variations dans les schémas de paiement modifient l'équilibre entre la compétition intragroupe et intergroupe dans chaque condition expérimentale :

- **Compétition inter + intragroupe (1)** : Les participants classés de la première à la cinquième place reçoivent chacun 15 €, tandis que les participants classés de la sixième à la dixième place reçoivent chacun 8 €.

- **Compétition inter + intragroupe (2)** : Les participants sont payés linéairement en fonction de leur classement. Le participant classé premier reçoit 15 €, celui classé deuxième 14 €, celui classé troisième 13 €, et ainsi de suite, jusqu'au dernier participant, qui reçoit 6 €.
- **Compétition inter + intragroupe (3)** : Le participant classé premier reçoit 20 €, ceux classés deuxième et troisième reçoivent 15 €, et les sept autres joueurs reçoivent chacun 10 €.

B.5.2 Performance et profils comportementaux des participants

L'analyse des différentes conditions expérimentales révèle des tendances claires quant aux comportements coopératifs et défecteurs selon le type de compétition en jeu.

Dans un premier temps, lorsque la compétition est absente et que les récompenses sont égales pour tous, les participants adoptent un comportement exploratoire, cherchant à découvrir de nouvelles cellules de haute valeur plutôt que de revisiter celles déjà identifiées. En revanche, dans les conditions impliquant de la compétition, des différences significatives émergent.

Table B.2 montre que la compétition intergroupe favorise les comportements coopératifs au sein des groupes, tandis que la compétition intragroupe encourage les comportements défecteurs. Ces conclusions s'alignent avec les résultats antérieurs, soulignant que la compétition intragroupe favorise la tromperie. Les conditions combinant les deux types de compétition présentent des niveaux de coopération variant entre ces deux extrêmes.

Ces observations corroborent les prédictions des modèles théoriques multi-échelles de l'évolution de la coopération, suggérant que lorsque la compétition entre groupes prévaut sur celle à l'intérieur des groupes, cela favorise la coopération intragroupe. Par exemple, dans une configuration où la compétition intergroupe est prédominante mais où il existe également une compétition intragroupe, une incitation forte à la collaboration est observée, puisque les individus peuvent être récompensés en étant classés parmi les cinq premiers.

Par ailleurs, même en l'absence de compétition formelle, l'introduction de récompenses différenciées associées à des performances entraîne une diminution du comportement coopératif. Cela suggère que la possibilité de pouvoir mesurer et quantifier ses performances par rapport à celles des autres encouragent la compétition, même sans structure compétitive explicite.

B.5.3 Modèle de prédiction des profils comportementaux individuels

Pour caractériser quantitativement les systèmes de paiement, trois observables sont introduits :

- (i) La moyenne des écarts-types des récompenses au sein des groupes, qui quantifie la compétition intragroupe.
- (ii) La différence de la récompense moyenne entre le groupe gagnant et le groupe perdant, qui quantifie la compétition intergroupe.

B.5. Impact de la Compétition Intragroupe et Intergroupe sur la Coopération Stigmergique

		INTERGROUP COMPETITION	
		NO	YES
INTRAGROUP COMPETITION	NO	No comp. + equal reward 84% – 13% – 3%	Intergroup comp. only 93% – 7% – 0%
		No comp. + performance reward 48% – 41% – 11%	
	YES	Intragroup comp. only 18% – 43% – 39%	Inter + intragroup comp. (1) 38% – 41% – 21%
			Inter + intragroup comp. (2) 32% – 42% – 26%
Inter + intragroup comp. (3) 13% – 46% – 41%			

TABLE B.2 : **Distribution des profils comportementaux.** Pourcentage de collaborateurs (vert), de neutres (marron) et de défecteurs (rouge) dans les sept conditions expérimentales, classées en fonction de la présence ou de l'absence de compétition intragroupe et intergroupe.




FIGURE B.7 : **Performance du modèle avec (i) et (ii).** Fractions prédites de collaborateurs, de neutres et de défecteurs en fonction des fractions expérimentales dans chaque expérience. La ligne en pointillé représente les prédictions parfaites.

(iii) La différence entre les récompenses moyennes des cinq joueurs les mieux classés et des cinq joueurs les moins bien classés, qui quantifie la compétition globale entre les dix participants et qui saisit certains éléments des compétitions intragroupe et intergroupe.

Pour mieux comprendre l'impact des structures de paiement sur les comportements individuels, nous avons construit un modèle linéaire qui intègre ces trois observables afin de prédire la proportion de chaque profil comportemental. Nous constatons que (i) est l'observable qui a de loin le meilleur pouvoir prédictif, puisqu'il permet déjà de différencier les sept conditions en fonction de la compétition intragroupe et intergroupe. Le deuxième observable le plus pertinent est (iii). Le modèle basé sur (i) et (iii) (voir Figure B.7) conduit alors à une excellente prédiction des fractions des trois profils de comportement dans chaque conditions. Enfin, nous constatons que l'observable (ii), qui donne une information sur la compétition intergroupe, ne présente qu'une faible corrélation avec les données. Cet observable est dominée par une seule condition (compétition intergroupe seulement) et n'est pas en mesure de faire la distinction entre les autres conditions. Finalement, le modèle linéaire incorporant les observables (i) et (iii) montre qu'une diminution de (i) ou une augmentation de (iii) est corrélé à une augmentation du nombre de collaborateurs et une diminution du nombre de défecteurs.

Notre analyse diffère de la perspective traditionnelle du domaine, qui catégorise souvent la compétition uniquement comme une combinaison de compétition intragroupe et intergroupe. Notre étude suggère que cette vision manichéenne ne permet pas, dans la pratique, de saisir l'interaction subtile entre les compétitions intragroupes et intergroupes, ainsi qu'une notion plus globale de la compétition indépendante des groupes.

B.6 Conclusion

Bien que les traces numériques jouent un rôle central dans les processus de prise de décision, leur impact sur ces derniers n'est pas encore totalement compris. Il est donc important d'étudier leur influence sur les comportements individuels et collectifs.

Pour participer à cette compréhension, nous avons développé un jeu dans lequel les participants s'engagent dans une tâche de recherche d'informations en utilisant un système de notation similaire à ceux que l'on trouve sur Internet. Ce jeu permet une analyse quantitative du comportement individuel et collectif, ainsi qu'une manipulation des effets de la compétition intragroupe et intergroupe sur la coopération. Bien qu'il ne capture pas toutes les complexités des systèmes de notation du monde réel, le jeu présente des similitudes et offre des options de personnalisation étendues, ce qui le rend adaptable à diverses conditions expérimentales.

En comparant deux expériences, l'une avec compétition entre les participants et l'autre sans, nous avons montré que des groupes peuvent utiliser efficacement des traces indirectes afin de coordonner leurs recherches d'information. Cependant, tout les individus n'adoptent pas le même comportement. Nous avons observé trois catégories distinctes : les collaborateurs, les neutres et les défecteurs, qui varient selon leur niveau de coopération. En absence de compétition, les individus présentent un comportement coopératif, générant des traces numériques fiables, tandis que la présence de compétition réduit la coopération, conduisant à des comportements plus neutres, voire trompeurs. Cette réduction de la coopération peut entraîner la rétention ou la diffusion d'informations erronées, diminuant ainsi la qualité des informations sociales sur le réseau.

Nous avons construit un modèle d'agent stochastique pour simuler les comportements humains observés dans le jeu. Ce modèle, divisé en deux parties : la stratégie de visite et la stratégie de notation, permet de modéliser de nombreux comportements différents grâce à sa paramétrisation. Sa simplicité permet une interprétation claire des mécanismes sous-jacents tout en restant ancrée dans les observations du monde réel. Le modèle permet d'améliorer la compréhension des comportements individuels et collectifs dans le jeu, de prédire le comportement humain dans divers scénarios et de contrôler le comportement des *bots* sociaux dans les expériences.

Par la suite, notre étude s'est penchée sur la manière dont les individus ajustent leurs stratégies en réponse à la dynamique de groupe et aux comportements des autres membres de leur groupe. Des expériences impliquant des interactions humaines avec des *bots* contrôlés par un modèle ont éclairé les comportements humains dans divers contextes de coopération. Nous avons observé une diversité de comportements chez les participants humains, influencés par le degré de coopération au sein du groupe, mettant en lumière leur capacité d'adaptation. Une analyse à l'aide d'un modèle linéaire a révélé que les participants évaluent le niveau de coopération du groupe à l'aide de signaux simples contenus dans les traces de couleurs et qu'ils ajustent leur comportement en conséquence. De surcroît, ces expériences ont souligné l'impact des *bots* sociaux sur le comportement humain, ouvrant la voie à la promotion de comportements coopératifs par le biais de manipulations délibérées.

Enfin, en variant les systèmes de paiement de nos expériences, nous avons étudié l'im-

l'impact de la compétition intragroupe et intergroupe sur les comportements et les performances individuels. Les résultats suggèrent que la compétition intergroupe encourage la coopération au sein des groupes, tandis que la compétition intragroupe favorise la tromperie. La combinaison des deux types de compétition aboutit à des niveaux intermédiaires de coopération, ce qui correspond aux modèles théoriques d'évolution de la coopération. Il est intéressant de noter que l'analyse révèle que le comportement individuel n'est pas uniquement influencé par la nature de la compétition. En effet, la capacité de mesurer et de comparer les performances entre les individus affecte également le comportement. Cela met en évidence la complexité des scénarios de compétition dans la vraie vie, dépassant la simple dynamique intragroupe et intergroupe.

Bibliography

- H. Ali Mehenni, S. Kobylyanskaya, I. Vasilescu, and L. Devillers. Nudges with conversational agents and social robots: A first experiment with children at a primary school. In L. F. D’Haro, Z. Callejas, and S. Nakamura, editors, *Conversational Dialogue Systems for the Next Decade*, volume 704, pages 257–270. Springer Singapore, Singapore, 2021.
- S. E. Asch. Opinions and social pressure. *Scientific American*, 193(5):31–35, 1955.
- R. Axelrod. *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*. Princeton University Press, 1997.
- R. M. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, NY, rev. ed edition, 2006.
- S. Balakrishnan and M. P. Koza. Information asymmetry, adverse selection and joint-ventures: Theory and evidence. *Journal of Economic Behavior & Organization*, 20(1):99–117, 1993.
- P. Ball. *Critical Mass: How One Thing Leads to Another*. Farrar, Straus and Giroux, New York, 1st american pbk. ed edition, 2006.
- P. Ball. *Why Society Is a Complex Matter*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.
- R. T. Ball and E. Ghysels. Automated earnings forecasts: Beat analysts or combine and conquer? *Management Science*, 64(10):4936–4952, 2018.
- R. K. Baltzersen. *Cultural-Historical Perspectives on Collective Intelligence*. Cambridge University Press, 2022.
- J. L. Barker, P. Barclay, and H. K. Reeve. Within-group competition reduces cooperation and payoffs in human groups. *Behavioral Ecology*, 23(4):735–741, 2012.
- A. Barrat and M. Weigt. On the properties of small-world network models. *The European Physical Journal B*, 13(3):547–560, 2000.
- A. Barrat, M. Barthelemy, and A. Vespignani. *Dynamical Processes on Complex Networks*. Cambridge University Press, Cambridge, UK ; New York, 2008.
- T. Bassanetti, S. Cezera, M. Delacroix, R. Escobedo, A. Blanchet, C. Sire, and G. Theraulaz. Cooperation and deception through stigmergic interactions in human groups. *Proceedings of the National Academy of Sciences*, 120(42), 2023.
- T. Bassanetti, S. Cezera, M. Delacroix, A. Blanchet, C. Sire, and G. Theraulaz. Quantifying the impact of social influence on collective intelligence in human groups with simple social bots [Unpublished manuscript]. 2024a.

- T. Bassanetti, S. Cezera, M. Delacroix, A. Blanchet, C. Sire, and G. Theraulaz. Impact of intragroup and intergroup competition on stigmergic cooperation [Unpublished manuscript]. 2024b.
- R. Beckers, J.-L. Deneubourg, S. Goss, and J. M. Pasteels. Collective decision making through food recruitment. *Insectes Sociaux*, 37(3):258–267, 1990.
- B. Beersma, J. R. Hollenbeck, S. E. Humphrey, H. Moon, D. E. Conlon, and D. R. Ilgen. Cooperation, competition, and team performance: Toward a contingency approach. *Academy of Management Journal*, 46(5):572–590, 2003.
- E. Bonabeau. Decisions 2.0: The power of collective intelligence. *MIT Sloan Management Review*, 50(2):45–52, 2009.
- T. Bose, A. Reina, and J. A. Marshall. Collective decision-making. *Current Opinion in Behavioral Sciences*, 16:30–34, 2017.
- E. Bothos, D. Apostolou, and G. Mentzas. Collective intelligence for idea management with Internet-based information aggregation markets. *Internet Research*, 19(1):26–41, 2009.
- E. Bothos, D. Apostolou, and G. Mentzas. Collective intelligence with web-based information aggregation markets: The role of market facilitation in idea management. *Expert Systems with Applications*, 39(1):1333–1345, 2012.
- D. A. Broniatowski, A. M. Jamison, S. H. Qi, L. AlKulaib, T. Chen, A. Benton, S. C. Quinn, and M. Dredze. Weaponized health communication: Twitter bots and Russian trolls amplify the vaccine debate. *American Journal of Public Health*, 108(10):1378–1384, 2018.
- J. L. Bronstein. *Mutualism*. Oxford university press, Oxford (GB), 2015.
- S. Camazine, J.-L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-Organization in Biological Systems*. Princeton University Press, 2001.
- D. Centola. *How Behavior Spreads: The Science of Complex Contagions*. Princeton university press, Princeton, N.J., 2018.
- D. Centola. The network science of collective intelligence. *Trends in Cognitive Sciences*, 26(11):923–941, 2022.
- D. Chavalarias. *Toxic Data*. Flammarion, Paris, 2022.
- R. Chiong and M. Kirley. The evolution of cooperation via stigmergic interactions. In *2012 IEEE Congress on Evolutionary Computation*, pages 1–8, 2012.
- R. B. Cialdini. *Influence: The Psychology of Persuasion*. HarperCollins, 2007.
- R. B. Cialdini and N. J. Goldstein. Social influence: Compliance and conformity. *Annual Review of Psychology*, 55(1974):591–621, 2004.

Bibliography

- M. Cinelli, G. de Francisci Morales, A. Galeazzi, W. Quattrociocchi, and M. Starnini. The echo chamber effect on social media. *Proceedings of the National Academy of Sciences*, 118(9):1–9, 2021.
- F. Cochoy, J. Hagberg, and M. McIntyre. *Digitalizing Consumption: How Devices Shape Consumer Culture*. Routledge, 2017.
- B. Collignon and C. Detrain. Self-organization and stigmergy. In C. K. Starr, editor, *Encyclopedia of Social Insects*, pages 793–803. Springer International Publishing, Cham, 2021.
- D. B. Cooney. The replicator dynamics for multilevel selection in evolutionary games. *Journal of Mathematical Biology*, 79(1):101–154, 2019.
- D. B. Cooney, S. A. Levin, Y. Mori, and J. B. Plotkin. Evolutionary dynamics within and among competing groups. *Proceedings of the National Academy of Sciences*, 120(20):e2216186120, 2023.
- I. Couzin. Collective minds. *Nature*, 445:715–715, 2007.
- É. Danchin, L.-A. Giraldeau, T. J. Valone, and R. H. Wagner. Public information: From nosy neighbors to cultural evolution. *Science*, 305(5683):487–491, 2004.
- A. C. Davison and D. V. Hinkley. *Bootstrap Methods and Their Application*. Cambridge University Press, 1997.
- N. de Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale, Paris, France, 1785.
- G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch. Mixing beliefs among interacting agents. *Advances in Complex Systems*, 03(01n04):87–98, 2000.
- J. L. Deneubourg, S. Aron, S. Goss, and J. M. Pasteels. The self-organizing exploratory pattern of the argentine ant. *Journal of Insect Behavior*, 3(2):159–168, 1990.
- M. J. Doyle and L. Marsh. Stigmergy 3.0: From ants to economies. *Cognitive Systems Research*, 21:1–6, 2013.
- J. Duboscq, V. Romano, A. MacIntosh, and C. Sueur. Social information transmission in animals: Lessons from studies of diffusion. *Frontiers in Psychology*, 7(AUG), 2016.
- L. A. Dugatkin. *Cooperation among Animals: A Evolutionary Perspective*. Oxford university press, New York Oxford, 1997.
- L. A. Dugatkin. Cooperation in animals: An evolutionary overview. *Biology & Philosophy*, 17(4):459–476, 2002.
- J. R. Dyer, C. C. Ioannou, L. J. Morrell, D. P. Croft, I. D. Couzin, D. A. Waters, and J. Krause. Consensus decision making in human crowds. *Animal Behaviour*, 75(2):461–470, 2008.

-
- B. Efron. Bootstrap methods: Another look at the jackknife. *The Annals of Statistics*, 7(1):1–26, 1979.
- E. Fehr and U. Fischbacher. The nature of human altruism. *Nature*, 425(6960):785–791, 2003.
- S. Feigin, G. Owens, and F. Goodyear-Smith. Theories of human altruism: A systematic review. *Journal of Psychiatry and Brain Functions*, 1(1), 2014.
- E. Ferrara, O. Varol, C. Davis, F. Menczer, and A. Flammini. The rise of social bots. *Communications of the ACM*, 59(7):96–104, 2016.
- J. Firth, J. Torous, B. Stubbs, J. A. Firth, G. Z. Steiner, L. Smith, M. Alvarez-Jimenez, J. Gleeson, D. Vancampfort, C. J. Armitage, and J. Sarris. The “online brain”: How the Internet may be changing our cognition. *World Psychiatry*, 18(2):119–129, 2019.
- P. Fraisse. De quelques comportements dits « spontanés ». *Enfance*, 21(3):161–181, 1968.
- N. E. Friedkin and E. C. Johnsen. Social influence networks and opinion change. *Advances in Group Processes*, 16:1–29, 1999.
- F. Galton. Vox populi. *Nature*, 75(1949):450–451, 1907.
- S. Garnier, J. Gautrais, and G. Theraulaz. The biological principles of swarm intelligence. *Swarm Intelligence*, 1(1):3–31, 2007.
- L. Gauvin, J. Vannimenus, and J.-P. Nadal. Phase diagram of a Schelling segregation model. *The European Physical Journal B*, 70(2):293–304, 2009.
- R. Gilbert. Looking for Mr. Schumpeter: Where are we in the competition–innovation debate? *Innovation Policy and the Economy*, 6:159–215, 2006.
- L. Giuggioli, J. R. Potts, D. I. Rubenstein, S. A. Levin, and R. M. May. Stigmergy, collective actions, and animal social spacing. *Proceedings of the National Academy of Sciences*, 110(42):16904–16909, 2013.
- E. S. Gloag, M. A. Javed, H. Wang, M. L. Gee, S. A. Wade, L. Turnbull, and C. B. Whitchurch. Stigmergy: A key driver of self-organization in bacterial biofilms. *Communicative and Integrative Biology*, 6(6), 2013.
- E. S. Gloag, L. Turnbull, and C. B. Whitchurch. Bacterial stigmergy: An organising principle of multicellular collective behaviours of bacteria. *Scientifica*, 2015:1–8, 2015.
- S. A. Golder and M. W. Macy. Digital footprints: Opportunities and challenges for online social research. *Annual Review of Sociology*, 40(1):129–152, 2014.
- R. L. Goldstone and T. M. Gureckis. Collective behavior. *Topics in Cognitive Science*, 1(3):412–438, 2009.

Bibliography

- R. L. Goldstone, M. E. Roberts, and T. M. Gureckis. Emergent processes in group behavior. *Current Directions in Psychological Science*, 17(1):10–15, 2008.
- J. A. Goncalo and M. M. Duguid. Follow the crowd in a new direction: When conformity pressure facilitates group creativity (and when it does not). *Organizational Behavior and Human Decision Processes*, 118(1):14–23, 2012.
- S. Goss, S. Aron, J. L. Deneubourg, and J. M. Pasteels. Self-organized shortcuts in the argentine ant. *Naturwissenschaften*, 76(12):579–581, 1989.
- U. Grandi and P. Turrini. A network-based rating system and its resistance to bribery. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, New York, NY, USA, 2016.
- M. S. Granovetter. The strength of weak ties. *American Journal of Sociology*, 78:1360–1380, 1973.
- 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 Sociaux*, 6(1):41–80, 1959.
- S. S. Gunasekaran, M. S. Ahmad, A. Tang, and S. A. Mostafa. The collective intelligence concept: A literature review from the behavioral and cognitive perspective. In *2016 2nd International Symposium on Agent, Multi-Agent Systems and Robotics (ISAMSR)*, pages 154–159, Bangi, Malaysia, 2016. IEEE.
- I. Gunes, C. Kaleli, A. Bilge, and H. Polat. Shilling attacks against recommender systems: A comprehensive survey. *Artificial Intelligence Review*, 42(4):767–799, 2014.
- D. Helbing, editor. *Social Self-Organization: Agent-based Simulations and Experiments to Study Emergent Social Behavior*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.
- D. Helbing, F. Schweitzer, J. Keltsch, and P. Molnár. Active walker model for the formation of human and animal trail systems. *Physical Review E*, 56(3):2527–2539, 1997.
- E. Hellinger. Neue begründung der theorie quadratischer formen von unendlichvielen veränderlichen. *Journal für die reine und angewandte Mathematik*, 1909(136):210–271, 1909.
- T. Hennig-Thurau, A. Marchand, and P. Marx. Can automated group recommender systems help consumers make better choices? *Journal of Marketing*, 76(5):89–109, 2012.
- J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Transaction on Information System*, 22(1):5–53, 2004.

-
- C. Heyes and B. Galef. *Social Learning in Animals: The Roots of Culture*. Academic Press, 1996.
- F. Heylighen. Stigmergy as a universal coordination mechanism I: Definition and components. *Cognitive Systems Research*, 38:4–13, 2016a.
- F. Heylighen. Stigmergy as a universal coordination mechanism II: Varieties and evolution. *Cognitive Systems Research*, 38:50–59, 2016b.
- L. Hong and S. E. Page. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences*, 101(46):16385–16389, 2004.
- N. Horsevad, D. Mateo, R. E. Kooij, A. Barrat, and R. Bouffanais. Transition from simple to complex contagion in collective decision-making. *Nature Communications*, 13(1):1442, 2022.
- B. Jayles, H. rin Kim, R. Escobedo, S. Cezera, A. Blanchet, T. Kameda, C. Sire, and G. Theraulaz. How social information can improve estimation accuracy in human groups. *Proceedings of the National Academy of Sciences*, 114(47):12620–12625, 2017.
- B. Jayles, R. Escobedo, S. Cezera, A. Blanchet, T. Kameda, C. Sire, and G. Theraulaz. The impact of incorrect social information on collective wisdom in human groups. *Journal of the Royal Society Interface*, 17(170), 2020a.
- B. Jayles, R. Escobedo, R. Pasqua, C. Zanon, A. Blanchet, M. Roy, G. Tredan, G. Theraulaz, and C. Sire. Collective information processing in human phase separation. *Philosophical Transactions of the Royal Society B*, 375(1807):20190801, 2020b.
- B. Jayles, C. Sire, and R. H. Kurvers. Crowd control: Reducing individual estimation bias by sharing biased social information. *PLoS Computational Biology*, 17(11 November):1–28, 2021a.
- B. Jayles, C. Sire, and R. H. J. M. Kurvers. Impact of sharing full versus averaged social information on social influence and estimation accuracy. *Journal of the Royal Society Interface*, 18(180):20210231, 2021b.
- M. Jesse and D. Jannach. Digital nudging with recommender systems: Survey and future directions. *Computers in Human Behavior Reports*, 3:100052, 2021.
- D. Kahneman, O. Sibony, and C. R. Sunstein. *Noise: A Flaw in Human Judgment*. Little, Brown Spark, New York, 2021.
- T. Kameda, T. Wisdom, W. Toyokawa, and K. Inukai. Is consensus-seeking unique to humans? A selective review of animal group decision-making and its implications for (human) social psychology. *Group Processes and Intergroup Relations*, 15(5):673–689, 2012.

Bibliography

- T. Kameda, W. Toyokawa, and R. S. Tindale. Information aggregation and collective intelligence beyond the wisdom of crowds. *Nature Reviews Psychology*, 0123456789, 2022.
- A. B. Kao and I. D. Couzin. Decision accuracy in complex environments is often maximized by small group sizes. *Proceedings of the Royal Society B*, 281(1784), 2014.
- A. B. Kao, A. M. Berdahl, A. T. Hartnett, M. J. Lutz, J. B. Bak-Coleman, C. C. Ioannou, X. Giam, and I. D. Couzin. Counteracting estimation bias and social influence to improve the wisdom of crowds. *Journal of the Royal Society Interface*, 15(141), 2018.
- S. Karanasios, D. Thakker, L. Lau, D. Allen, V. Dimitrova, and A. Norman. Making sense of digital traces: An activity theory driven ontological approach. *Journal of the American Society for Information Science and Technology*, 64(12):2452–2467, 2013.
- A. Khuong, J. Gautrais, A. Perna, C. Sbaï, M. Combe, P. Kuntz, C. Jost, and G. Theraulaz. Stigmergic construction and topochemical information shape ant nest architecture. *Proceedings of the National Academy of Sciences*, 113(5):1303–1308, 2016.
- A. J. King and D. J. Sumpter. Murmurations. *Current Biology*, 22(4):R112–R114, 2012.
- P. Kitcher. The evolution of human altruism. *The Journal of Philosophy*, 90(10):497–516, 1993.
- A. Kittur, B. Lee, and R. E. Kraut. Coordination in collective intelligence: The role of team structure and task interdependence. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1495–1504, Boston MA USA, 2009. ACM.
- M. Kusters and J. Van der Heijden. From mechanism to virtue: Evaluating nudge theory. *Evaluation*, 21(3):276–291, 2015.
- B. Kramer and A. MacKinnon. Localization: Theory and experiment. *Reports on Progress in Physics*, 56(12):1469–1564, 1993.
- R. H. Kurvers, U. Hertz, J. Karpus, M. P. Balode, B. Jayles, K. Binmore, and B. Bahrami. Strategic disinformation outperforms honesty in competition for social influence. *iScience*, 24(12):103505, 2021.
- K. N. Laland. Social learning strategies. *Animal Learning & Behavior*, 32(1):4–14, 2004.
- Y.-H. Li and X. Tian. Quorum sensing and bacterial social interactions in biofilms. *Sensors*, 12(3):2519–2538, 2012.
- K. K. Loh and R. Kanai. How has the internet reshaped human cognition? *Neuroscientist*, 22(5):506–520, 2016.
- J. Lorenz, H. Rauhut, F. Schweitzer, and D. Helbing. How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences*, 108(22):9020–9025, 2011.

-
- J. Loveland and J. Reagle. Wikipedia and encyclopedic production. *New Media & Society*, 15(8):1294–1311, 2013.
- L. Lü, M. Medo, C. H. Yeung, Y.-C. Zhang, Z.-K. Zhang, and T. Zhou. Recommender systems. *Physics Reports*, 519(1):1–49, 2012.
- S. Luo. A unifying framework reveals key properties of multilevel selection. *Journal of Theoretical Biology*, 341:41–52, 2014.
- T. W. Malone and M. S. Bernstein, editors. *Handbook of Collective Intelligence*. The MIT Press, Cambridge, Massachusetts London, England, 2015.
- J. A. Marshall, R. H. J. M. Kurvers, J. Krause, and M. Wolf. Quorums enable optimal pooling of independent judgements in biological systems. *eLife*, 8:1–14, 2019.
- P. Mavrodiev, C. J. Tessone, and F. Schweitzer. Quantifying the effects of social influence. *Scientific Reports*, 3, 2013.
- C. Mele, T. Russo Spena, V. Kaartemo, and M. L. Marzullo. Smart nudging: How cognitive technologies enable choice architectures for value co-creation. *Journal of Business Research*, 129:949–960, 2021.
- A. P. Melis and D. Semmann. How is human cooperation different? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1553):2663–2674, 2010.
- S. Milgram, L. Bickman, and L. Berkowitz. Note on the drawing power of crowds of different size. *Journal of Personality and Social Psychology*, 13(2):79–82, 1969.
- M. Mokkonen and C. Lindstedt. The evolutionary ecology of deception. *Biological Reviews*, 91(4):1020–1035, 2016.
- S. Moscovici. *Psychologie sociale*. PUF, Paris, nouvelle éd. edition, 2014.
- A. Moseley and G. Stoker. Nudging citizens? Prospects and pitfalls confronting a new heuristic. *Resources, Conservation and Recycling*, 79:4–10, 2013.
- M. Moussaid, S. Garnier, G. Theraulaz, and D. Helbing. Collective information processing and pattern formation in swarms, flocks, and crowds. *Topics in Cognitive Science*, 1(3): 469–497, 2009.
- M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45(2): 167–256, 2003.
- G. Nicolis and C. Nicolis. *Foundations of Complex Systems: Emergence, Information and Prediction*. World Scientific, 2 edition, 2012.
- M. A. Nowak. Five rules for the evolution of cooperation. *Science*, 314(5805):1560–1563, 2006.

Bibliography

- C. O’Neil. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Broadway Books, New York, first paperback edition, 2017.
- S. E. Page. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton Univ. Press, Princeton, NJ, 2008.
- K. Panchanathan and R. Boyd. Indirect reciprocity can stabilize cooperation without the second-order free rider problem. *Nature*, 432(7016):499–502, 2004.
- M. Pereda, V. Capraro, and A. Sánchez. Group size effects and critical mass in public goods games. *Scientific Reports*, 9(1):5503, 2019.
- D. Pollard. *A User’s Guide to Measure Theoretic Probability*. Cambridge University Press, 1 edition, 2001.
- D. G. Rand and M. A. Nowak. Human cooperation. *Trends in Cognitive Sciences*, 17(8): 413–425, 2013.
- A. Rezgui and K. Crowston. Stigmergic coordination in Wikipedia. In *Proceedings of the 14th International Symposium on Open Collaboration*, pages 1–12, Paris France, 2018. ACM.
- P. J. Richerson, R. J. Baldini, A. V. Bell, K. Demps, K. Frost, V. Hillis, S. Mathew, E. K. Newton, N. Naar, L. Newson, C. Ross, P. E. Smaldino, T. M. Waring, and M. Zefferman. Cultural group selection plays an essential role in explaining human cooperation: A sketch of the evidence. *Behavioral and Brain Sciences*, 39(2016), 2016.
- J. Salminen. Collective intelligence in humans: A literature review, 2012.
- C. D. Santos, S. Przybyzin, M. Wikelski, and D. K. Dechmann. Collective decision-making in homing pigeons: Larger flocks take longer to decide but do not make better decisions. *PLoS ONE*, 11(2):1–8, 2016.
- F. P. Santos, J. M. Pacheco, and F. C. Santos. The complexity of human cooperation under indirect reciprocity. *Philosophical Transactions of the Royal Society B*, 2021.
- M. Schöbel, J. Rieskamp, and R. Huber. Social influences in sequential decision making. *PLoS ONE*, 11(1):1–23, 2016.
- T. D. Seeley. *Honeybee Democracy*. Princeton University Press, 2010.
- B. Simon, J. A. Fletcher, and M. Doebeli. Towards a general theory of group selection. *Evolution*, 67(6):1561–1572, 2013.
- L. Sjöberg. Are all crowds equally wise? A comparison of political election. *Journal of Forecasting*, 18(September 2008):1–18, 2009.
- A. P. Smith. An investigation of the mechanisms underlying nest construction in the mud wasp *Paralastor* sp. (Hymenoptera: Eumenidae). *Animal Behaviour*, 26:232–240, 1978.

- T. A. Society and C. Press. Competition versus cooperation : Success of individuals foraging alone and in groups. *The American Naturalist*, 142(1):42–58, 1993.
- W. Steinel and C. K. De Dreu. Social motives and strategic misrepresentation in social decision making. *Journal of Personality and Social Psychology*, 86(3):419–434, 2004.
- W. Steinel, S. Utz, and 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. *Organizational Behavior and Human Decision Processes*, 113(2):85–96, 2010.
- T. Stonier. *Information Technology, Collective Intelligence and the Evolution of Human Societies*, pages 85–106. Springer London, London, 1992.
- D. J. T. Sumpter. *Collective Animal Behavior*. Princeton University Press, 2010.
- J. Surowiecki. *The Wisdom of Crowds*. Anchor, 2005.
- R. H. Thaler and C. R. Sunstein. *Nudge: Improving Decisions about Health, Wealth, and Happiness*. Yale university press, New Haven (Conn.), 2008.
- G. Theraulaz and E. Bonabeau. Coordination in distributed building. *Science*, 269(5224):686–688, 1995.
- G. Theraulaz and E. Bonabeau. A brief history of stigmergy. *Artificial Life*, 5(2):97–116, 1999.
- B. Thierry, G. Theraulaz, J. Y. Gautier, and B. Stiegler. Joint memory. *Behavioural Processes*, 35:127–140, 1996.
- C. Toma and F. Butera. Cooperation versus competition effects on information sharing and use in group decision-making. *Social and Personality Psychology Compass*, 9(9):455–467, 2015.
- A. Traulsen and M. A. Nowak. Evolution of cooperation by multilevel selection. *Proceedings of the National Academy of Sciences*, 103(29):10952–10955, 2006.
- A. Tversky and D. Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.
- H. Van Dyke Parunak. A survey of environments and mechanisms for human-human stigmergy. In D. Weyns, H. Van Dyke Parunak, and F. Michel, editors, *Environments for Multi-Agent Systems II*, volume 3830, pages 163–186, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- D. J. Watts and S. H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684):440–442, 1998.
- D. S. Wilson. A theory of group selection. *Proceedings of the National Academy of Sciences*, 72(1):143–146, 1975.

Bibliography

- D. S. Wilson. *Does Altruism Exist? Culture, Genes, and the Welfare of Others*. Yale University Press : Templeton Press, New Haven ; London, 2015.
- D. S. Wilson and L. A. Dugatkin. Group selection and assortative interactions. *The American Naturalist*, 149(2):336–351, 1997.
- A. W. Woolley, I. Aggarwal, and T. W. Malone. Collective intelligence and group performance. *Current Directions in Psychological Science*, 24(6):420–424, 2015.
- D. E. Zand. Trust and managerial problem solving. *Administrative Science Quarterly*, 17(2):229–239, 1972.
- T. R. Zentall and B. G. Galef Jr. *Social Learning: Psychological and Biological Perspectives*. Lawrence Erlbaum Associates, Inc, 1988.

Titre : Impact de différents schémas de compétition sur les processus de coopération stigmergiques au sein de groupes humains

Mots clés : intelligence collective, stigmergie, influence sociale, traces numériques, modélisation et simulation numérique

Résumé : La stigmergie constitue un mécanisme de coordination générique largement exploité dans les sociétés animales, dans lequel les traces laissées par un individu dans l'environnement guident et stimulent les actions ultérieures de cet individu et d'autres individus. Dans le contexte humain, avec la numérisation de la société, de nouvelles formes de processus stigmergiques ont émergé à travers le développement de services en ligne qui exploitent largement les traces numériques laissées par leurs utilisateurs, notamment par le biais de systèmes de recommandation ou d'évaluation. Dans ce contexte, il est essentiel de comprendre l'influence de ces traces numériques sur la prise de décision individuelle et collective. Dans un premier temps, j'analyse et modélise l'interaction de groupes d'individus avec leurs traces numériques, et détermine comment ils peuvent exploiter ces traces pour coopérer dans une tâche de recherche d'informations. Par la suite, j'étudie l'influence de la compétition, tant au sein du groupe qu'entre les groupes, sur la manière dont les individus coopèrent pour mener à bien cette tâche. Pour répondre à ces questions, nous avons conçu le jeu multijoueur en ligne Stigmer, que nous avons utilisé pour mener 16 séries d'expériences. Dans ce jeu novateur, des groupes d'individus déposent et exploitent des traces numériques dans le cadre d'une tâche de recherche d'informations qui intègre un système d'évaluation à 5-étoiles. Ce système est similaire à ceux utilisés par de nombreuses plateformes en ligne, où les utilisateurs peuvent évaluer des produits, des services ou des vendeurs. Dans le jeu, les participants interagissent sur le même tableau de valeurs cachées, recherchant les cellules ayant les valeurs les plus élevées, en exploitant uniquement les informations indirectes fournies sous la forme de traces colorées résultant de leurs évaluations collectives. Ce cadre expérimental contrôlé permet une analyse quantitative approfondie du comportement individuel et collectif, en offrant la possibilité de manipuler et d'étudier l'impact combiné de la compétition intra et intergroupe sur la dynamique de coopération. Les résultats expérimentaux et de la modélisation montrent que le type et l'intensité de la compétition influent sur la manière dont les individus interprètent et utilisent les traces numériques, et sur la fiabilité des informations fournies par ces traces. Cette étude révèle que les individus peuvent être classés en trois profils comportementaux qui diffèrent par leur degré de coopération : les collaborateurs, les neutres et les trompeurs. En l'absence de compétition, les traces numériques induisent spontanément une coopération entre les individus, soulignant ainsi le potentiel des processus stigmergiques pour favoriser la collaboration dans les groupes humains. De même, la compétition entre deux groupes favorise le comportement coopératif des membres d'un groupe qui cherchent à surpasser la performance des membres de l'autre groupe. Cependant, la compétition au sein d'un groupe peut engendrer des comportements trompeurs, où les individus manipulent leurs évaluations pour obtenir un avantage compétitif sur les autres membres du groupe. Ainsi, dans les processus de prise de décisions, une information sociale non fiable renforce l'utilisation d'informations privées au détriment de cette information sociale. Enfin, les situations qui combinent à la fois de la compétition intragroupe et intergroupe font apparaître des niveaux variables de coopération entre les individus, expliqués par notre étude. En mettant en lumière les liens entre compétition, coopération, tromperie, et prise de décision, ces travaux établissent les fondements pour comprendre des interactions stigmergiques dans les environnements numériques. Ces résultats peuvent contribuer au développement d'algorithmes de prise de décision personnalisés et de systèmes d'intelligence collective artificielle fondés sur la stigmergie.

Title: Impact of different competition schemes on stigmergic cooperation processes in human groups

Key words: collective intelligence, stigmergy, social influence, digital traces, computational modeling

Abstract: Stigmergy is a generic coordination mechanism widely used by animal societies, in which traces left by individuals in the environment guide and stimulate the subsequent actions of the same or different individuals. In the human context, with the digitization of society, new forms of stigmergic processes have emerged through the development of online services that extensively exploit the digital traces left by their users, in particular, using rating-based recommendation systems. Therefore, understanding the impact of these digital traces on both individual and collective decision-making is essential. This study pursues two main objectives. First, I investigate and modelize the interactions of groups of individuals with their digital traces, and determine how they can exploit these traces to cooperate in an information search task. Subsequently, the research explores the impact of intragroup and intergroup competition on the dynamics of cooperation in the framework of this information search task. To answer these questions, we have developed the online multiplayer Stigmer game, on which we base 16 series of experiments under varying conditions. In this game, groups of individuals leave and exploit digital traces in an information search task that implements a 5-star rating system. This system is similar to recommendation systems used by many online marketplaces and platforms, where users can evaluate products, services, or sellers. In the game, all individuals interact with a grid of hidden values, searching for cells with the highest values, and using only indirect information provided in the form of colored traces resulting from their collective ratings. This controlled environment allows for a thorough and quantitative analysis of individual and collective behaviors, and offers the possibility of manipulating and studying the combined impact of intragroup and intergroup competition on cooperation. The experimental and modeling results indicate that the type and intensity of competition determine how individuals interpret and use digital traces, and impact the reliability of the information delivered via these traces. This study reveals that individuals can be classified into three behavioral profiles that differ in their degree of cooperation: collaborators, neutrals, and defectors. When there is no competition, digital traces spontaneously induce cooperation among individuals, highlighting the potential for stigmergic processes to foster collaboration in human groups. Likewise, competition between two groups also promotes cooperative behavior among group members who aim to outperform the members of the other group. However, intragroup competition can prompt deceptive behaviors, as individuals may manipulate their ratings to gain a competitive advantage over the other group members. In this situation, the presence of misinformation reinforces the use of private information over social information in the decision-making process. Finally, situations that combine both intragroup and intergroup competition display varying levels of cooperation between individuals, that we explain. This research establishes the foundations for understanding stigmergic interactions in digital environments, shedding light on the relationships between competition, cooperation, deception, and decision-making. The insights gained may contribute to the development of sustainable and cooperative personalized decision-making algorithms and artificial collective intelligence systems grounded in stigmergy.