nature portfolio

Corresponding author(s): Igor Kagan

Last updated by author(s): YYYY-MM-DD

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our Editorial Policies and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed	
☐	☒	The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒	A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☒	☐	A description of all covariates tested
☐	☒	A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒	A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒	For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☒	☐	For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐	For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒	Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	Data collection was conducted using a custom Python 3.10 code, available at https://gitlab.gwdg.de/darius.lewen/lewen_et_al, and linked to at https://osf.io/56hw7/
Data analysis	Data analysis was conducted using a custom Python 3.10 code and R version 4.4.2, available at https://gitlab.gwdg.de/darius.lewen/lewen_et_al, and linked to at https://osf.io/56hw7/

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our policy

All data collected in the experiment and used for the analyses is available at https://osf.io/56hw7/ and https://gitlab.gwdg.de/darius.lewen/lewen_et_al

Research involving human participants, their data, or biological material

Policy information about studies with human participants or human data. See also policy information about sex, gender (identity/presentation), and sexual orientation and race, ethnicity and racism.

Reporting on sex and gender	The analyzed dataset included 116 participants, of whom 19 self-identified as female and the remainder as male. The predominance of male participants was due to recruitment overlap with a related ongoing study investigating hormonal effects in males. These participants responded affirmatively to the question "Bist du chromosomal geschlechtlich männlich?" ("Are you chromosomally male?"). No analyses based on sex or gender were conducted; therefore, the terms "sex" and "gender" are not used throughout the manuscript. Self-reported gender information is provided in the openly available dataset.
Reporting on race, ethnicity, or other socially relevant groupings	We did not collect or report data on participants' race, ethnicity, or other socially defined groupings, as these variables were not pertinent to the research questions addressed in this study.
Population characteristics	Please see above and below.
Recruitment	Participants were recruited via student chat groups, ads, and flyers around the Goettingen campus, and were predominantly Bachelor, Master, or PhD German and international students. As such, the sample is biased toward individuals with a typical educational and socio-economic background for university students, likely characterized by higher education levels, middle to upper socio-economic status, and familiarity with academic settings. This may potentially limit the generalizability of the findings to broader or more diverse populations.
Ethics oversight	The experimental protocol was approved by the ethics committee of the Georg-Elias-Mueller-Institute for Psychology, University of Goettingen.

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☐ Life sciences ☒ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	Quantitative experimental and computational modeling study, using continuous 2D cursor movement data from pairs of participants collecting cooperative and competitive targets in a foraging game in a real-time face-to-face environment.
Research sample	Predominantly Bachelor, Master or PhD German and international students, age range 18–36 years, 116 participants in 58 dyads (19 females, 97 males, mean ± SD age: 25 ± 4 years, range 18–36 years, 46 male dyads, 7 female dyads, 5 mixed female/male dyads). Most participants were male because they were recruited for a related ongoing study of male hormonal effects.
Sampling strategy	A convenience sampling strategy was employed to recruit approximately 60 dyads. As this was a novel experimental paradigm with no established effect sizes or outcome distributions, the target sample size was informed by previous research on social interaction and sensorimotor decision-making. The sample size was selected to balance practical feasibility with the goal of capturing a potentially broad range of emergent dyadic behaviors and interaction strategies.
Data collection	The experiments were conducted in the Dyadic Interaction Platform laboratory at the German Primate Center, as detailed in the Methods section. In brief, two participants sat face-to-face and played a visually guided Cooperation–Competition Foraging game. Each participant used a computer mouse to control movements of a cursor ("agent") within a 2D field on a large, bidirectional transparent OLED screen positioned between participants. An experimenter—either one of the study authors or a trained student assistant—initiated and supervised each session from a control console behind a partition, remaining out of direct view of the participants. The authors were fully informed about the study design and general goals and had real-time visual access to the interaction via a mirrored display. However, they had no means to influence participants' behavior and were blind to the cumulative dynamics of dyadic strategies, which required offline analysis. The only real-time information available to them during the task was the ongoing accumulation of payoffs.
Timing	Data were collected between January 2022 and December 2022.
Data exclusions	As stated in the Methods, 4 out of 62 dyads were excluded from the current analysis because they exhibited large differences in behavior between blocks. The exclusion criterion was not pre-established but was developed during the initial analysis, with the goal to focus on stable behavioral patterns.

Non-participation	No participants dropped out.
Randomization	Participants were not allocated into experimental groups.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study	
☒	☐	Antibodies
☒	☐	Eukaryotic cell lines
☒	☐	Palaeontology and archaeology
☒	☐	Animals and other organisms
☒	☐	Clinical data
☒	☐	Dual use research of concern
☒	☐	Plants

Methods

n/a	Involved in the study	
☒	☐	ChIP-seq
☒	☐	Flow cytometry
☒	☐	MRI-based neuroimaging

Plants

Seed stocks	N/A
Novel plant genotypes	N/A
Authentication	N/A

nature portfolio

Peer Review File

Continuous dynamics of cooperation and competition in social decision-making

Corresponding Author: Dr Igor Kagan

This file contains all editorial decision letters in order by version, followed by all author rebuttals in order by version.

Version 0:

Decision Letter:

** Please ensure you delete the link to your author homepage in this e-mail if you wish to forward it to your coauthors **

Dear Dr Kagan,

Thank you for your patience during the peer-review process. Your manuscript titled "Continuous dynamics of cooperation and competition in social foraging" has now been seen by 3 reviewers, and I include their comments at the end of this message. They find your work of interest but raised some important points. We are interested in the possibility of publishing your study in Communications Psychology, but would like to consider your responses to these concerns and assess a revised manuscript before we make a final decision on publication.

We therefore invite you to revise and resubmit your manuscript, along with a point-by-point response to the reviewers. Please highlight all changes in the manuscript text file.

Editorially, we ask you to carefully address the methodological and conceptual concerns raised by R1 and 2 (e.g., details about how skill of players to play the game was measured). Please also consider re-organizing the results part to improve the clarity, as suggested by R3.

We are committed to providing a fair and constructive peer-review process. Please don't hesitate to contact us if you wish to discuss the revision in more detail.

I am attaching an Editorial Requests Table that details critical reporting requirements for the revised manuscript. Please attend to each item and ensure your manuscript is fully compliant. If your revised manuscript is not aligned with these requests on major issues, such as those concerning statistics, it may be returned to you for further revisions without re-review.

Please submit the following items:

- Revised manuscript
- Point-by-point response to the referees' comments
- Cover letter (as a separate document)
- Nature Research Reporting Summary
- Completed Editorial Request Table (attached).

via this link: Link Redacted .

** This url links to your confidential home page and associated information about manuscripts you may have submitted or be reviewing for us. If you wish to forward this email to co-authors, please delete the link to your homepage first **

Additional guidance is available in our style and formatting guide Communications Psychology formatting guide.

We hope to receive your revised paper within 8 weeks; please let us know if you aren't able to submit it within this time so that we can discuss how best to proceed. If we don't hear from you, and the revision process takes significantly longer, we may close your file. In this event, we will still be happy to reconsider your paper at a later date, provided it still presents a significant contribution to the literature at that stage.

We would appreciate it if you could keep us informed about an estimated timescale for resubmission, to facilitate our planning.

We look forward to seeing the revised manuscript and thank you for the opportunity to review your work.

Best regards,

Troby Lui, on behalf of

Yafeng Pan

Troby Lui, PhD
Associate Editor
Communications Psychology

Yafeng Pan, PhD
Editorial Board Member
Communications Psychology
orcid.org/0000-0002-5633-8313

REVIEWER EXPERTISE:
Reviewer #1: cooperation and competition, social decision-making, computational modeling
Reviewer #2: cooperation and competition, social decision-making, computational modeling
Reviewer #3: cooperation and competition, social decision-making, computational modeling

REVIEWER REPORTS:

Reviewer #1 (Remarks to the Author):

The manuscript Continuous dynamics of cooperation and competition in social foraging considered for submission in Communications psychology introduces a novel experimental paradigm based on a game played in pairs facing each others and interacting via a transparent screen separating the two players. The task consists in developing strategies, either competing with the other player or cooperating with them, to get financial rewards in a spatially-explicit continuous environment. The manuscript presents the paradigm, an experiment made of 2x20 min trials from 62 different pairs and reports a detailed and thorough analysis of the collected data.

Overall, the manuscript is very well written, figures (including captions) of high quality and the provided analysis exhaustive and in-depth. The reported results support well the motivations of studying this novel experimental paradigm further. I particularly appreciated that in addition to the novel experimental paradigm, many in-depth ideas about how to analyse the data coming from the experiments are already provided.

My main concern regards how skill of players to play the game is measured. I am not fully convinced that what has been measured by the authors actually reflects skills of players and improves the understanding of the system. Skill in the manuscript has been defined as the difference in single targets collected, normalized by the total number of single targets. This choice was motivated by the finding that the correlation between payoff difference and single target difference was very high. This correlation is however heavily biased since trivially perfect for pairs with high FSTs, which, by definition, are only discriminated by single target differences. Also, by construction, this definition excludes pairs that chose to avoid the single targets in their strategies. Therefore, this metrics seems biased in its explanatory power in favour of high FSTs, cannot encompass all the strategies exhibited by players and may not be fully linked to skills of players. It seems, for instance, that strategies themselves may lead the players to be in very different mindsets – highly coordinated strategies in which players know with high certainty how the other will play (either focusing on joint targets or dividing the arena in two) presumably leading players to play more automatically with a smaller cognitive load (as suggested by the authors L553-554). In that case, the difference seen in payoff may be more a disinterest in the actual final payoff than clear skill differences.

I am also sceptical about the reference to foraging theory, especially in the title of the manuscript: it is indeed interesting to draw connections with behavioural ecology and animal decision-making and to discuss those but I am not sure the game primarily intends to investigate the questions associated with these fields. This paradigm surely allows to investigate human decision-making in the context of a game with small impacts on the lives of players but I find the connection to animal foraging to be a fairly bold claim.

Minor comments:

L192: There is no subpanel in Figure S2, S2a is not required

Fig 3d, caption: maybe adding (b) to the sentence "The actual dyads lie between simulated strategies" would make clearer that the two lines are the same ones as on panel b.

Fig 4a, caption: what is 'invite placements'? This is not explained yet when we first get the figure introduced (L212) – explanation only comes L231.

L388 Although the statement of this conclusive paragraph (that cooperative dryads get lower joint payoffs) is in general true on Fig 6a, it could be slightly nuanced by what the authors reported above: when cooperation is highly coordinated (diamonds), they result in higher payoffs compared to dyads with similar FSTs and the actual best dyad overall could also be described as cooperation (the blue + symbol) since the two players divided the space equally and actively avoid competition.

L651: rolled a die > a dice

Reviewer #2 (Remarks to the Author):

In the manuscript entitled "Continuous dynamics of cooperation and competition in social foraging" the authors propose a new game, the Cooperation-Competition-Foraging game, to study social interactions of dyads navigating a continuous shared space while foraging for rewards. Members of the dyads can collect rewards either together or independently.

The task involves a free-flowing social game, where two participants use continuous arm movements to collect targets in a face-to-face setting with full transparency about their co-player movements and appearance. The task consists of moving a mouse on a 2D screen towards either one of two cooperative targets to be reached together with the other player or towards an individual target to be reached before the other player.

Subjects played for two blocks of 20 minutes each and after the first 10-15 minutes most dyads' behaviours conformed to stable strategies with groups of dyads opting for either mainly cooperative, mainly competitive or intermediate strategies as estimated by the fraction of single targets reached over a short period of interactions.

The authors derive optimal strategies based on principles of path minimisations and cooperative strategies to achieve it and show that a theoretically optimal fraction of single targets (FST) stands between 0.33 and 0.4 depending on whether participants move together or spread themselves in the space following single targets.

They develop a model based on these principles and on observed behavioural patterns of prioritising targets not chosen on the previous trial and prioritising cooperation after competition. They further characterise classes of behaviour and the spatio-temporal features underlying them as well as the impact of difference in skills in the game.

Overall this is a really nice paper, building on a fantastic new task to study the trade off between cooperation and competition, a multifaceted, engaging and dynamic social foraging game on 2D that taps in the continuous nature of social interactions as well as the sensorimotor behaviour that underlies them. The game provides for an extremely rich repertoire of social behaviours which is both well modelled and well described by a compelling set of analysis and a commendable collection of videos. I congratulate the authors for their work and I am really glad if they found any inspiration on our earlier work on the Space Dilemma.

My concerns are minor and are mostly aimed at increasing clarity and understanding for the reader.

My most significant clarification relates to the definition of the most advantageous position. How the authors derive the advantageous position in equation 11 and 12 is theoretically sound but I found the colormap in figure 3c is somewhat confusing. I understand the map should identify the most advantageous position to minimise the distance to the next target irrespective of whether this is the single target or the joint target closest to the collecting agent. So it follows that, intuitively, it should be the combination of two maps, one relating to the set of best positions to get to the closest joint target, and one to the position that minimises the dyad distance to the new individual target, each weighted by the relative target type weighting parameter. The former map would be a circus centered around the joint target with radius equal to the distance from the agent who collected the previous round (like in sup fig.1b). The latter map would look something like sup fig1a, with the best position being symmetrical with respect to the midpoint from the collecting agent. However in fig.3c it seems like the gradient mainly reflects the distance from the closest joint target. Can you provide an intuition both in the legend and text of why that is instead of a blob around the advantageous position that combines the two maps? Is it because in that example the weight is highly biased towards joint collection? In general, the weight will make a difference on what this advantageous position heatmap look like and it would be good if that was articulated in the text and perhaps with two different example heatmaps for different weights.

Relatedly in figure S1 b, it's counterintuitive why position closer to the closest joint target should be preferable to positions equidistant from the closest joint but further away from the collecting agent (as symmetrical from the midpoint as possible). My understanding is that if only the maximum distance to the closest joint matters there is no point for the non collecting agent in being closer to it than the collecting agent. Could the authors perhaps show the two maps representing the maximum distance to the closest joint target and the minimum average distance to the unknown individual target both separately and combined? This would potentially make those maps more intuitive. At the moment I can't get my head round why S1 b is qualitatively different than S1c

In figure 3a the weights for the distance are not quite clear. Are they represented visually somewhere? From the legend it seems only the weighted distances are. So is the orange weighted distance always higher than the blue one because the original distances are? That's not immediately apparent from the figure. The "joint preferred" condition with three different weighted distances can be confusing. Could you also somehow visually or numerically represent the weights? Or report the numbers in the legend? That would help understand more intuitively why on the left joint wins and on the right the individual target. Perhaps also highlight the winning target in each condition.

Likewise in figure 3b it would also help to get a numerical sense of the three weights conditions, in the legend or in the panel.

Can the authors swap panel 3b and 3c? Conceptually it makes sense to explain graphically the advantageous position before the simulation and current panel b and d could look nicer side by side?

Caption figure 3d most dyads don't seem to be in between simulated strategies? Not clear what the red lines measure. Is it the distance from the "always same starting position" curve (That is, what reveals their degree of adv starting position)? It doesn't look so from the plot. Why?

Have the authors plan to modulate behaviour by having different rewards structures/conditions in future experiment and can the authors comment in the discussion what happens if one eliminates the collaborative targets? It seems to me that would provide a 2D generalisation of the intermediate condition in the Space Dilemma.

Can the authors comment on whether the reduction in the shift towards the advantageous position depends on the risk of increasing the chance of the co-player winning the next single trial and therefore can be framed as a matter of trust?

Minor points

The diamond dyads do the turn-take which is interesting. From the video it looks like they do it almost normatively, irrespective of distance of the other two targets. Is that so or there are exceptions for the situations where those other targets are closer?

Even though the sample was highly biased, did you observe any gender differences?

Reviewer #3 (Remarks to the Author):

This study marks a novel approach in the fields of both game theory and foraging. They use a "transparent game" method to study continuous dynamics of cooperation and competition in a foraging setting. I really like their approach and find it to be an important contribution to multiple fields, in terms of both the unique findings and the methodological approach. However, I have some suggestions for improvement regarding readability of the paper:

1. The Introduction is comprehensive but it is quite lengthy and redundant at times. I would suggest making it more concise. It also focuses a lot on the methodological advancements, which is warranted, but it seems to be the only focus. In a way I was unprepared for some elements of the task design (such as single and joint targets) and of the results ("invitation", different strategies, skill difference etc.). I think the authors can foreshadow these concepts in the introduction and motivate these aspects of their design, in addition to current focus on continuous, face-to-face dynamics.

2. I found that the results are also quite verbose and hard to follow at some times because of the high amount of information. Maybe the authors can move some plots/sub-results to the supplementary to make it easier for a reader to grasp the results. It also would be helpful to end the introduction or start the results by explaining the different strategies and what are their signatures.

3. Line 128 - "transporting" reads weird.

4. In one of the results subheadings, there is "shaping" -- the authors probably mean "shapes/shape".

* TRANSPARENT PEER REVIEW: Communications Psychology uses a transparent peer review system. This means that we publish the editorial decision letters including Reviewers' comments to the authors and the author rebuttal letters online as a supplementary peer review file. However, on author request, confidential information and data can be removed from the published reviewer reports and rebuttal letters prior to publication. If your manuscript has been previously reviewed at another journal, those Reviewers' comments would not form part of the published peer review file.

** Visit Nature Research's author and referees' website at www.nature.com/authors for information about policies, services and author benefits**

Communications Psychology is committed to improving transparency in authorship. As part of our efforts in this direction, we are now requesting that all authors identified as ‘corresponding author’ create and link their Open Researcher and Contributor Identifier (ORCID) with their account on the Manuscript Tracking System prior to acceptance. ORCID helps the scientific community achieve unambiguous attribution of all scholarly contributions. You can create and link your ORCID from the home page of the Manuscript Tracking System by clicking on ‘Modify my Springer Nature account’ and following the instructions in the link below. Please also inform all co-authors that they can add their ORCIDs to their accounts and that they must do so prior to acceptance.
https://www.springernature.com/gp/researchers/orcid/orcid-for-nature-research

For more information please visit http://www.springernature.com/orcid

If you experience problems in linking your ORCID, please contact the Platform Support Helpdesk.

Version 1:

Decision Letter:

** Please ensure you delete the link to your author homepage in this e-mail if you wish to forward it to your coauthors **

Dear Dr Kagan,

Your manuscript titled "Continuous dynamics of cooperation and competition in social decision-making" has now been seen by our reviewers, whose comments appear below. In light of their advice I am delighted to say that we are happy, in principle, to publish a suitably revised version in Communications Psychology.

We therefore invite you to revise your paper one last time to address the remaining concerns of our reviewers and a list of editorial requests. At the same time we ask that you edit your manuscript to comply with our format requirements and to maximise the accessibility and therefore the impact of your work.

EDITORIAL REQUESTS:

Please review our specific editorial comments and requests regarding your manuscript in the attached "Editorial Requests Table". Please outline your response to each request in the right hand column. Please upload the completed table with your manuscript files as a Related Manuscript file.

If you have any questions or concerns about any of our requests, please do not hesitate to contact me.

SUBMISSION INFORMATION:

In order to accept your paper, we require the files listed here https://www.nature.com/documents/commsj-file-checklist.pdf .

OPEN ACCESS:

Communications Psychology is a fully open access journal. Articles are made freely accessible on publication. For further information about article processing charges, open access funding, and advice and support from Nature Research, please visit https://www.nature.com/commspsychol/open-access

At acceptance, you will be provided with instructions for completing the open access licence agreement on behalf of all authors. This grants us the necessary permissions to publish your paper. Additionally, you will be asked to declare that all required third party permissions have been obtained, and to provide billing information in order to pay the article-processing charge (APC).

* TRANSPARENT PEER REVIEW: Communications Psychology uses a transparent peer review system. On author request, confidential information and data can be removed from the published reviewer reports and rebuttal letters prior to publication. If you are concerned about the release of confidential data, please let us know specifically what information you would like to have removed. Please note that we cannot incorporate redactions for any other reasons.

* CODE AVAILABILITY: All Communications Psychology manuscripts must include a section titled "Code Availability" at the end of the methods section. We require that the custom analysis code supporting your conclusions is made available in a publicly accessible repository at this stage; please choose a repository that generates a digital object identifier (DOI) for the code; the link to the repository and the DOI must be included in the Code Availability statement. Publication as Supplementary Information will not suffice.

* DATA AVAILABILITY:
All Communications Psychology manuscripts must include a section titled "Data Availability" at the end of the Methods section. More information on this policy, is available in the Editorial Requests Table and at http://www.nature.com/authors/policies/data/data-availability-statements-data-citations.pdf.

Please use the following link to submit the above items:
Link Redacted
** This url links to your confidential home page and associated information about manuscripts you may have submitted or be reviewing for us. If you wish to forward this email to co-authors, please delete the link to your homepage first **

We hope to hear from you within two weeks; please let us know if you need more time.

Best regards,

Troby Lui, on behalf of

Yafeng Pan

Troby Lui, PhD
Associate Editor
Communications Psychology

Yafeng Pan, PhD
Editorial Board Member
Communications Psychology
orcid.org/0000-0002-5633-8313

REVIEWERS' COMMENTS:

Reviewer #1 (Remarks to the Author):

The authors appropriately addressed all my comments.

I spotted one typo:
L398: ‘hightest’ > highest

Reviewer #2 (Remarks to the Author):

I am happy with the authors' revisions and for the paper to be published.

Reviewer #3 (Remarks to the Author):

I thank the authors for their revisions and for this work! I have no further comments.

MAX PLANCK INSTITUTE
FOR DYNAMICS AND
SELF-ORGANIZATION

DPZ
Deutsches Primatenzentrum
Leibniz-Institut für Primatenforschung

Prof. Viola Priesemann
Max Planck Institute for Dynamics and Self-Organization
Am Fassberg 17, 37077, Göttingen
viola.priesemann@ds.mpg.de

Dr. Igor Kagan
German Primate Center – Leibniz Institute for Primate Research
Kellnerweg 4, 37077, Göttingen
ikagan@dpz.eu

July 17, 2025

Troby Lui, PhD
Associate Editor

Yafeng Pan, PhD
Editorial Board Member

Response to Reviewers

Dear Dr. Lui, dear Dr. Pan, dear Reviewers,

We are grateful for the opportunity to submit the revised manuscript, which incorporates the revisions requested by the three Reviewers. We would like to express our appreciation for the very fast and positive evaluation of our work, and for thoughtful and constructive questions and suggestions that significantly improved the manuscript.

To summarize our revisions, we carefully addressed the methodological and conceptual questions raised by the Reviewers, such as providing more details about skill metric, and clarifying advantageous placement and target type weighting. We also made the Introduction more concise, improved a lead-up to key concepts, and streamlined the Results to improve the clarity.

Below we provide a point-by-point reply to the Reviewers' comments. The comments are in *italics*; our responses are in the regular font.

In the revised article file with tracked changes, the substantial changes are highlighted in light blue. We refrained from copying the revised text to this response for brevity, and use the line numbers (L...) to indicate the location of specific changes in response to the comments.

Reviewer 1

[...] Overall, the manuscript is very well written, figures (including captions) of high quality and the provided analysis exhaustive and in-depth. The reported results support well the motivations of studying this novel experimental paradigm further. I particularly appreciated that in addition to the novel experimental paradigm, many in-depth ideas about how to analyse the data coming from the experiments are already provided.

We thank Reviewer 1 for the positive evaluation and the appreciation of our analysis approaches.

My main concern regards how skill of players to play the game is measured. I am not fully convinced that what has been measured by the authors actually reflects skills of players and improves the understanding of the system. Skill in the manuscript has been defined as the difference in single targets collected, normalized by the total number of single targets. This choice was motivated by the finding that the correlation between payoff difference and single target difference was very high. This correlation is however heavily biased since trivially perfect for pairs with high FSTs, which, by definition, are only discriminated by single target differences. Also, by construction, this definition excludes pairs that chose to avoid the single targets in their strategies. Therefore, this metrics seems biased in its explanatory power in favour of high FSTs, cannot encompass all the strategies exhibited by players and may not be fully linked to skills of players. It seems, for instance, that strategies themselves may lead the players to be in very different mindsets – highly coordinated strategies in which players know with high certainty how the other will play (either focusing on joint targets or dividing the arena in two) presumably leading players to play more automatically with a smaller cognitive load (as suggested by the authors L553-554). In that case, the difference seen in payoff may be more a disinterest in the actual final payoff than clear skill differences.

We agree with the Reviewer's concern and appreciate the opportunity to clarify and improve this aspect of our manuscript. First, we note that the measure of skill based on single target collection was calculated only for collection cycles in which **both** participants actively aimed for single targets (as we now clearly state in the Results, L412-420, and the revised Fig. 7 caption, as well as in the Methods, L759-765). We acknowledge that this was not emphasized clearly enough in the Results and the figure caption, and our previous explanation in the Methods section was confusing, which might have contributed to a misunderstanding of this metric.

Second, even if imperfect, this approach provides a way to quantify an important aspect of social behavior — namely, the cost of cooperation, particularly for the more skilled player who could have, counterfactually, adhered to a more competitive strategy to maximize individual payoff. We contend that regardless of the underlying factors driving less successful collection of single target by one of the players — be it disinterest in the payoff or actual sensorimotor ability —- the measure still reflects a relative advantage, motivational or sensorimotor, of the better-performing player in competitive contexts.

That said, we fully recognize that this measure of skill is limited: its reliability is inherently dependent on the frequency of single target collections, it cannot be estimated in dyads who avoided single targets entirely, and it does not capture cooperative skill in dyads with stable cooperative tactics (e.g., field division or alternating turns). Additionally, as the reviewer points out, strategic choice — particularly in highly coordinated dyads — can reduce cognitive load and make performance less reflective of individual competitive sensorimotor abilities. These factors limit the generality of our current metric.

We have revised the manuscript to make these limitations explicit. We now introduce it as "**competitive skill**", and explain the underlying assumptions: i) if both agents move straight to the single target, they compete, and ii) the competitive skill is independent of FST (L412-420, and L437-438). In the Limitations section of the Discussion (L619-626), we now describe this measure as a post-hoc and potentially biased proxy for competitive skill. We also emphasize that a more robust and generalizable assessment might require an independent skill calibration block, in which participants are explicitly instructed to compete for single targets under controlled conditions. Such a manipulation would allow for a cleaner dissociation between strategic choices and individual motor or planning abilities. While our current data do not permit this, we present the analysis as a useful step toward understanding the relationship between competitive skill differences, payoff inequality, and cost of cooperation, within certain strategic regimes.

I am also sceptical about the reference to foraging theory, especially in the title of the manuscript: it is indeed interesting to draw connections with behavioural ecology and animal decision-making and to discuss those but I am not sure the game primarily intends to investigate the questions associated with these fields. This paradigm surely allows to investigate human decision-making in the context of a game with small impacts on the lives of players but I find the connection to animal foraging to be a fairly bold claim.

We appreciate the reviewer's critique and agree that our paradigm differs from naturalistic animal foraging in terms of ecological stakes. We have changed the title to "Continuous dynamics of cooperation and competition in social **decision-making**" to emphasize the decision-making framework while maintaining the key features of dynamic interaction.

Nonetheless, we respectfully argue that core principles from foraging theory remain highly relevant. Even in laboratory settings with modest rewards, the foraging framework provides a well-established and generalizable approach to studying cost-benefit trade-offs, spatiotemporal choices, and strategy adaptation under uncertainty — all of which are present in our task. Indeed, a substantial body of human and primate neuroscience research has successfully applied foraging theory to analyze stay/leave decisions, reward history effects, information sampling, and exploration–exploitation dynamics under similar low-stakes conditions (e.g. [1-6]). Likewise, our game instantiates a continuous decision space in which participants must weigh action costs, opportunity costs, and the predicted behavior of a social partner across a sequence of choices —- hallmarks of foraging-like decisions. We have added a clarification to the Discussion to articulate this reasoning more clearly (L543-551).

[1] Kolling, N., Behrens, T. E. J., Mars, R. B., Rushworth, M. F. S. (2012). Neural mechanisms of foraging. Science, 336(6077), 95–98.

[2] Hayden, B. Y., Pearson, J. M., Platt, M. L. (2011). Neuronal basis of sequential foraging decisions in a patchy environment. Nature Neuroscience, 14(7), 933–939.

[3] Pearson, J. M., Hayden, B. Y., Raghavachari, S., Platt, M. L. (2009). Neurons in posterior cingulate cortex signal exploratory decisions in a dynamic multioption choice task. Current Biology, 19(18), 1532–1537.

[4] Hall-McMaster, S., Luyckx, F. (2019). Revisiting foraging approaches in neuroscience. Cognitive, Affective, Behavioral Neuroscience, 19(2), 225–230.

[5] Gabay, A. S., Apps, M. A. J. (2021). Foraging optimally in social neuroscience: computations and methodological considerations. Social Cognitive and Affective Neuroscience, 16(8), 782–794.

[6] Mobbs, D., Trimmer, P. C., Blumstein, D. T., Dayan, P. (2018). Foraging for foundations in decision neuroscience: insights from ethology. Nature Reviews Neuroscience, 19(7), 419–427.

Minor comments:

L192: There is no subpanel in Figure S2, S2a is not required

Thank you for noticing the inconsistency, done.

Fig 3d, caption: maybe adding (b) to the sentence "The actual dyads lie between simulated strategies" would make clearer that the two lines are the same ones as on panel b.

Thank you, we extensively revised this caption (also in response to Reviewer 2), please see page 7.

Fig 4a, caption: what is 'invite placements'? This is not explained yet when we first get the figure introduced (L212) – explanation only comes L231.

We placed this figure after the text where invitations are first introduced.

L388 Although the statement of this conclusive paragraph (that cooperative dryads get lower joint payoffs) is in general true on Fig 6a, it could be slightly nuanced by what the authors reported above: when cooperation is highly coordinated (diamonds), they result in higher payoffs compared to dyads with similar FSTs and the actual best dyad overall could also be described as cooperation (the blue + symbol) since the two players divided the space equally and actively avoid competition.

Thank you, we added these points (L393-401).

L651: rolled a die > a dice

Respectfully, die is singular, dice is plural (the form "rolled the dice" might be more familiar because dice are frequently used in pairs in games).

Reviewer 2

[...] Overall this is a really nice paper, building on a fantastic new task to study the trade off between cooperation and competition, a multifaceted, engaging and dynamic social foraging game on 2D that taps in the continuous nature of social interactions as well as the sensorimotor behaviour that underlies them. The game provides for an extremely rich repertoire of social behaviours which is both well modelled and well described by a compelling set of analysis and a commendable collection of videos. I congratulate the authors for their work and I am really glad if they found any inspiration on our earlier work on the Space Dilemma.

We thank Reviewer 2 for very positive and in-depth evaluation, and for their earlier work on the

Space Dilemma that indeed shaped much of our thinking and interpretations.

My concerns are minor and are mostly aimed at increasing clarity and understanding for the reader.

My most significant clarification relates to the definition of the most advantageous position. How the authors derive the advantageous position in equation 11 and 12 is theoretically sound but I found the colormap in figure 3c is somewhat confusing. I understand the map should identify the most advantageous position to minimise the distance to the next target irrespective of whether this is the single target or the joint target closest to the collecting agent. So it follows that, intuitively, it should be the combination of two maps, one relating to the set of best positions to get to the closest joint target, and one to the position that minimises the dyad distance to the new individual target, each weighted by the relative target type weighting parameter. The former map would be a circus centered around the joint target with radius equal to the distance from the agent who collected the previous round (like in sup fig.1b). The latter map would look something like sup fig1a, with the best position being symmetrical with respect to the midpoint from the collecting agent. However in fig.3c it seems like the gradient mainly reflects the distance from the closest joint target. Can you provide an intuition both in the legend and text of why that is instead of a blob around the advantageous position that combines the two maps? Is it because in that example the weight is highly biased towards joint collection? In general, the weight will make a difference on what this advantageous position heatmap look like and it would be good if that was articulated in the text and perhaps with two different example heatmaps for different weights.

Relatedly in figure S1b, it's counterintuitive why position closer to the closest joint target should be preferable to positions equidistant from the closest joint but further away from the collecting agent (as symmetrical from the midpoint as possible). My understanding is that if only the maximum distance to the closest joint matters there is no point for the non collecting agent in being closer to it than the collecting agent. Could the authors perhaps show the two maps representing the maximum distance to the closest joint target and the minimum average distance to the unknown individual target both separately and combined? This would potentially make those maps more intuitive. At the moment I can't get my head round why S1b is qualitatively different than S1c.

We appreciate the thoughtful and detailed nature of these inquiries, which highlighted a lack of clarity in our previous presentation, for which we apologize. The reviewer's intuition and understanding are correct. The panel in the original Fig. 3c (now 3b) — the same panel reproduced in Fig. S1c — is already a weighted combination of the two maps: (i) the principle demonstrated in Fig. S1a ("ignore joint targets; minimize distance to single target"), and (ii) the principle minimizing the distance to the closest joint target. Both panels S1b and S1c represent the combination of these principles, but the relative contribution of single and joint targets depends on the distance to the closest joint target, and the single/joint target weighting. Due to the fixed color scale and saturation of the "joint target map" in S1b, the contribution of the single target map (as in S1a) was not very apparent.

We now replotted these panels with a more intuitively oriented and better scaled color bar that clearly indicates the cutoff below 6.5 cm, added explanations and arrows pointing to the optimal placement(s) to each panel, and extensively revised Figure S1 and Fig. 3 captions. We also explicitly indicate the weighting factor (w=0.5 in Fig. S1 panels (a-d), w=0.99 in (e)). We also updated the Results (L165-166).

In figure 3a the weights for the distance are not quite clear. Are they represented visually somewhere?

From the legend it seems only the weighted distances are. So is the orange weighted distance always higher than the blue one because the original distances are? That's not immediately apparent from the figure. The "joint preferred" condition with three different weighted distances can be confusing. Could you also somehow visually or numerically represent the weights? Or report the numbers in the legend? That would help understand more intuitively why on the left joint wins and on the right the individual target. Perhaps also highlight the winning target in each condition. Likewise in figure 3b it would also help to get a numerical sense of the three weights conditions, in the legend or in the panel.

Good points; we added the actual weights and also indicate which target is collected with an asterisk. Yes, in the examples shown, the orange weighted distance is always higher than the blue one because of the original distances. We now mention that the weights in previous panel 3b (now 3c) are the same as in 3a.

Can the authors swap panel 3b and 3c? Conceptually it makes sense to explain graphically the advantageous position before the simulation and current panel b and d could look nicer side by side?

Great suggestion, we swapped the panels.

Caption figure 3d most dyads don't seem to be in between simulated strategies? Not clear what the red lines measure. Is it the distance from the "always same starting position" curve (That is, what reveals their degree of adv starting position)? It doesn't look so from the plot. Why?

Yes, the red lines indicate the distance reduction due to some degree of advantageous placement, i.e. the distance reduction due to the free agent not sharing the same starting positions with the collecting agent, as in the simulations that yield the "always same stating position" curves. For the FST=1 dyads, the tops of the red lines do not perfectly align with the dark gray dotted curve because the simulations reflect an idealized approximation over a very large number of target collections, while the dyads' actual means are based on finite, variable data and therefore include some jitter. We clarified these points in the legend in Fig. 3d, and in the extensively revised caption.

Have the authors plan to modulate behaviour by having different rewards structures/conditions in future experiment and can the authors comment in the discussion what happens if one eliminates the collaborative targets? It seems to me that would provide a 2D generalization of the intermediate condition in the Space Dilemma.

Yes, we plan to build on this game to investigate different aspects of human and macaque behavior, including manipulations of payoff structure — similarly to what has been done in the Space Dilemma study in Pisauro et al. 2022 — and information availability. But our currently ongoing project uses the same formulation in a round-robin design, investigating the relationship between personality and hormonal traits and convergence to a specific strategy, and neural correlates of those strategies. Reviewer 2 is absolutely correct that the CCF game without the joint targets would be a generalization of the Space Dilemma in 2D, adding the target collection movements and mutually observed real-time actions and thus, continuity not only in space but also in time.

Can the authors comment on whether the reduction in the shift towards the advantageous position depends on the risk of increasing the chance of the co-player winning the next single trial and therefore can be framed as a matter of trust?

Yes, the (non-competitive) advantageous placement can indeed increase the co-player's chance of winning the next single target, and thus can be framed as a potential indicator of trust that other one will do the same. However, interpreting it as a direct measure of trust is challenging. In an idealized scenario —- with no skill difference and instantaneous reaction times —- the competitive placement would be right next to the collecting agent, and any extra distance towards the center could be attributed to trust. In practice, however, this strategy might not be very obvious to the naive participants, and real players might try to gain a competitive **timing** advantage by shifting toward the center. Therefore, this positioning can reflect either trust (non-competitive, advantageous placement for either player) or competitive heuristics based on perceived skill and reaction times, making it difficult to disentangle the two without incorporating actual movement skill and reaction times into a model.

We note that the invitations to the joint targets constitute a more reliable measure of trust. We now mention both of these aspects in the Discussion (L505 and the next paragraph, L506-513).

Minor points

The diamond dyads do the turn-take which is interesting. From the video it looks like they do it almost normatively, irrespective of distance of the other two targets. Is that so or there are exceptions for the situations where those other targets are closer?

Indeed, these dyads ignored the distance and performed strict normative turn-taking — we now emphasize it more in L279-282.

Even though the sample was highly biased, did you observe any gender differences?

It is undoubtedly very interesting question but beyond the scope of our study due to the sample limitations, as appreciated by the Reviewer. Anecdotally, few female and mixed dyads we tested (7 female dyads, 5 mixed female/male dyads) exhibited a similar distribution of FST values, from 0 to 1, showing cooperative, intermediate and competitive strategies.

Reviewer 3

This study marks a novel approach in the fields of both game theory and foraging. They use a "transparent game" method to study continuous dynamics of cooperation and competition in a foraging setting. I really like their approach and find it to be an important contribution to multiple fields, in terms of both the unique findings and the methodological approach. However, I have some suggestions for improvement regarding readability of the paper:

We thank Reviewer 3 for the positive feedback and for recognizing the important contributions of our work. We appreciate these suggestions and have revised the manuscript to improve its readability.

1. The Introduction is comprehensive but it is quite lengthy and redundant at times. I would suggest making it more concise. It also focuses a lot on the methodological advancements, which is warranted, but it seems to be the only focus. In a way I was unprepared for some elements of the task design (such as single and joint targets) and of the results ("invitation", different strategies, skill difference etc.). I think the authors can foreshadow these concepts in the introduction and motivate these aspects of their *design, in addition to current focus on continuous, face-to-face dynamics.*

We made the Introduction more concise but also expanded the lead-up to the key design concepts (L72-82).

2. I found that the results are also quite verbose and hard to follow at some times because of the high amount of information. Maybe the authors can move some plots/sub-results to the supplementary to make it easier for a reader to grasp the results. It also would be helpful to end the introduction or start the results by explaining the different strategies and what are their signatures.

We systematically streamlined the narrative in the Results, and added an overview of strategies right after we introduce the three main groups (L139-142). We retained most figures because we felt strongly that they are necessary to present a logical story, but we moved the panels (d), (e) and (f) from Fig. 6, and the corresponding text, to the new Suppl. Fig. S5.

3. Line 128 - "transporting" reads weird.

Thank you, we changed the sentence to: "By embedding foraging in a shared virtual space and a salient social context, our setup provides a controlled yet dynamic environment to study how dyads develop cooperative or competitive strategies."

4. In one of the results subheadings, there is "shaping" – the authors probably mean "shapes/shape".

We changed it to "shape".

Once again, we thank all Reviewers for their time and effort!

Sincerely,

Darius Lewen, Viola Priesemann and Igor Kagan, on behalf of all authors

Description of Additional Supplementary Files

File name- Supplementary Movie S1

File description – Setup and game demonstration. Human Dyadic Interaction Platform setup and the game demonstration (60 s), followed by the replay of an example intermediate dyad.

File name- Supplementary Movie S2

File description – Cooperative example. Replay of a representative dyad from the cooperative group.

File name- Supplementary Movie S3

File description – Intermediate strategy example. Replay of a representative dyad from the intermediate group.

File name- Supplementary Movie S4
File description – Competition example. Replay of a representative dyad from the competitive group.

File name- Supplementary Movie S5
File description – Invitations examples. Collection cycles where one agent invites the other to a joint target.

File name- Supplementary Movie S6
File description – Cooperative turn-taking. Replay of one of three dyads who alternated between the two joint targets.

File name- Supplementary Movie S7

File description – Strongly curved trajectories. Examples of collection cycles featuring strongly curved trajectories, reflecting initial miscoordination and changes of mind.

File name- Supplementary Movie S8

File description – Competitive placement example. Replay of a dyad performing competitive advantageous placement

File name- Supplementary Movie S9
File description – Cooperative placement for single targets. A special dyad that achieved nearly optimal advantageous placement by splitting the game field.

Supplementary Information

Continuous dynamics of cooperation and competition in social decision-making

Darius Lewen[1,2†], Vladyslav Ivanov[3†], Jonas Dehning[1,2†], Johannes Ruß[2], Anna Fischer[2], Lars Penke[2], Anne Schacht[2], Alexander Gail[2,3], Viola Priesemann[1,2‡*], Igor Kagan[3‡*]

[1] Max Planck Institute for Dynamics and Self-Organization, Am Faßberg 17, Göttingen, Germany.
[2] University of Göttingen, Wilhelmsplatz 1, Göttingen, Germany.
[3] German Primate Center – Leibniz Institute for Primate Research, Kellnerweg 4, Göttingen, Germany.

† These authors contributed equally.
‡ These authors jointly supervised this work.

* Corresponding authors: viola.priesemann@ds.mpg.de, ikagan@dpz.eu

This PDF file includes:

Supplementary text 1. Supplementary Methods, including Supplementary Tables S1 to S5
Supplementary Figures S1 to S6
Supplementary Movies S1 to S9
Supplementary text 2. Instructions for participants

1 Supplementary Methods

1.1 Optimal dyad strategy

Table S1 | Variables used in section 1.1, and their descriptions.

Variable	Description
A, B	Agent names.
$\vec{a}_i, \vec{b}_i$	Agent positions at the start of cycle i.
$\vec{s}_i$	Single target position at the start of cycle i.
$\vec{j}_i^A$	Predominantly blue joint target (higher payoff share for agent A) position at the start of cycle i.
$\vec{j}_i^B$	Predominantly orange joint target (higher payoff share for agent B) position at the start of cycle i.
D_S	Relevant distance for the single target.
D_J	Relevant distance for a joint target.
D	Shortest relevant distance.
w	Target-type weighting parameter, weights single vs joint targets.
$\vec{t}_{i+1}$	Target position updating rule.
$\vec{x}_{i+1}$	Agent position updating rule.
$\vec{y}_{i+1}$	Extended agent position updating rule with advantageous placement.
N	Total number of collection cycles.
Φ	Fraction of single targets collected.

To establish an analytical foundation, we derive the optimal strategy for a dyad that maximizes their joint payoff. We assume for simplicity agents with identical speed, making the Euclidean distances from agents' positions $\vec{a}$ and $\vec{b}$ to the targets $\vec{s}, \vec{j}^A, \vec{j}^B$ the sole determinant of the payoff per second for each target choice. Given that the collection process of a single target begins when the first agent reaches it, the relevant distance for single targets D_S is the minimum distance between the agents and the target:

$$D_S(\vec{a}, \vec{b}, \vec{s}) = \min(||\vec{a} - \vec{s}||, ||\vec{b} - \vec{s}||). \tag{1}$$

Conversely, for joint targets, the collection process starts when the last agent arrives, so the relevant distance D_J is the maximum distance between the agents and the target:

$$D_J(\vec{a}, \vec{b}, \vec{j}^X) = \max(||\vec{a} - \vec{j}^X||, ||\vec{b} - \vec{j}^X||). \tag{2}$$

To maximize the joint payoff, we assume the dyad always selects the target with the shortest relevant distance D:

$$D(\vec{a}, \vec{b}, \vec{s}, \vec{j}^A, \vec{j}^B, w) = \min \begin{cases} D_S(\vec{a}, \vec{b}, \vec{s}) \cdot (1-w) \\ D_J(\vec{a}, \vec{b}, \vec{j}^A) \cdot w \\ D_J(\vec{a}, \vec{b}, \vec{j}^B) \cdot w \end{cases}, \tag{3}$$

where w is a target-type weighting parameter. When $w = 1/2$, both target types are equally weighted, approximating the optimal dyad strategy. For $w \neq 1/2$, we obtain a non-optimal strategy that approximates the best strategy among those collecting the same fraction of single targets (FST value).
To simulate this strategy, we define the initial agent positions at the center of the game field:

$$\vec{a}_0 = \vec{b}_0 = \begin{pmatrix} 1/2 \\ 1/2, \end{pmatrix} \tag{4}$$

and sample the initial positions of the three targets from the two-dimensional standard uniform distribution:

$$\vec{s}_0, \vec{j}_0^A, \vec{j}_0^B \sim \mathcal{U}^2_{[0,1]}. \tag{5}$$

We define the target position updating rule $\vec{t}_{i+1}$ as implemented in the game

$$\vec{t}_{i+1} = \begin{cases} \vec{u} \sim \mathcal{U}^2_{[0,1]} & \text{if } \vec{t}_i \text{ was collected} \\ \vec{t}_i & \text{otherwise.} \end{cases} \tag{6}$$

If a target is collected, its new position is sampled from the two-dimensional standard uniform distribution $\mathcal{U}^2_{[0,1]}$; otherwise, it remains unchanged. In the first analysis of the optimal strategy, we assume both agents share the same position at the beginning of a collection cycle, defining the agent position updating rule $\vec{x}_{i+1}$ such that they are always at the previously collected target's position at the start of each collection cycle:

$$\vec{x}_{i+1} = \begin{cases} \vec{s}_i & \text{if } \vec{s}_i \text{ was collected} \\ \vec{j}_i^A & \text{if } \vec{j}_i^A \text{ was collected} \\ \vec{j}_i^B & \text{otherwise.} \end{cases} \tag{7}$$

With these updating rules, we can simulate each dyad strategy and thus calculate the expected fraction of single targets Φ:

$$\mathbb{E}[\Phi|w] = \lim_{N\to\infty} \frac{1}{N} \sum_{i=0}^{N} \begin{cases} 1 & \text{if } \vec{s}_i \text{ was collected} \\ 0 & \text{otherwise.} \end{cases} \tag{8}$$

and the expected relevant distance:

$$\mathbb{E}[D|w] = \lim_{N\to\infty} \frac{1}{N} \sum_{i=0}^{N} D(\vec{a}_i, \vec{b}_i, \vec{s}_i, \vec{j}_i^A, \vec{j}_i^B, w), \tag{9}$$

which is proportional to the payoff per second (see the gray curve in Fig. 3c for simulation results).

To relax the initial assumption of agents always sharing the same position at the start of each collection cycle, we extend the agent position updating rule: after collecting a joint target, both agents inevitably share a similar position, therefore, the updating rule remains in this case unchanged. However, during a single target collection, the non-collecting agent can position itself advantageously for the next collection cycle. Thus, we introduce an extended agent position updating rule $\vec{y}_{i+1}$:

$$\vec{y}_{i+1} = \begin{cases} \hat{\vec{y}}_{i+1} & \text{if } \vec{s}_i \text{ was collected by the other agent, and} \\ \vec{x}_{i+1} & \text{otherwise.} \end{cases} \tag{10}$$

This advantageous placement $\hat{\vec{y}}_{i+1}$ minimizes the expected relevant distance:

$$\hat{\vec{y}}_{i+1} = \arg\min_{\vec{y}} \mathbb{E}[D|\vec{x}_{i+1}, \vec{y}, \vec{j}_i^A, \vec{j}_i^B, w]. \tag{11}$$

Note here that $\vec{x}_{i+1} = \vec{s}_i$ since the previous collection was a single target collection. The expected relevant distance equals the integral over all possible single target spawn positions:

$$\mathbb{E}[D|\vec{x}, \vec{y}, \vec{j}^A, \vec{j}^B, w] = \int_{[0,1]^2} D(\vec{x}, \vec{y}, \vec{s}, \vec{j}^A, \vec{j}^B, w) d\vec{s}. \tag{12}$$

Fig. 3c (black curve) illustrates the simulation results with advantageous placement, and Fig. 3d and Supplementary Fig. S1 depict the expected relevant distances for various free-agent placements.
By simulating strategies for various $w \in [0, 1]$, we establish a function that maps target choices to FST values, $w \mapsto \Phi$ (Supplementary Fig. S2). Inverting this function allows us to predict dyads' target choices based on their FST value Φ and the current game configuration $\vec{a}_i, \vec{b}_i, \vec{s}_i, \vec{j}_i^A, \vec{j}_i^B$. The subsequent section will delve deeper into modeling dyad strategies.

1.2 Modeling dyad strategies

Table S2 | Variables used in section 1.2, and their descriptions.

Variable	Description
j	Index of collection type. Possible values: 1 $\Leftrightarrow$ single target collection, 2 $\Leftrightarrow$ predominantly blue joint target collection, and 3 $\Leftrightarrow$ predominantly orange joint target collection.
$\hat{y}_j^{(i)}$	Predicted probability for target collection j in collection cycle i.
Softmax	Softmax function.
$\boldsymbol{\theta}$	Regression coefficient matrix.
$\vec{x}^{(i)}$	Covariates of collection cycle i.
$\vec{x}_d$	Covariate component containing the relevant distances for the single and joint targets. The collection cycle i is omitted for simplicity.
$\vec{x}_p$	Covariate component encoding the identity of the previously collected target, whether an invite is present, and towards which joint target the invite is directed. The collection cycle i is omitted for simplicity.
$\vec{x}_{p'}$	Covariate component encoding the identity of the target collected before the previous one. The collection cycle i is omitted for simplicity.
$\boldsymbol{\theta}_d$	Component of the regression coefficient matrix that weights the relevant distance of the single target against those of the joint targets.
d_0	Regression coefficient of the single target's relevant distance.
d_1	Regression coefficient of the joint targets' relevant distances.
$\boldsymbol{\theta}_p$	Component of the regression coefficient matrix encoding either (1) if there is no invite, the relative change in collection probability of the previously collected target, or (2) if an invite is present, the relative change in collection probability of the joint target towards which the invitation is directed.
p_0	Regression coefficient for the change in the collection probability of the single target relative to others if the previous target was a single target.
p_1	Regression coefficient for the change in the collection probability of a joint target relative to others if the previous target was this joint target.
p_2	Regression coefficient for the change in the collection probability of a joint target relative to others if an invite towards this joint target is present.
$\boldsymbol{\theta}_{p'}$	Component of the regression coefficient matrix encoding the relative change in collection probability of the target collected before the previously collected target.
p_3	Regression coefficient for the relative change in the collection probability of the target collected before the previous one. This coefficient accounts for both single and joint targets.
$\mathcal{L}$	Negative log-likelihood function of the generalized linear model (GLM). It is the Cross-Entropy loss.
N	Total number of collection cycles.
$y_j^{(i)}$	Binary encoding indicating whether target collection j occurred in collection cycle i.

For each individual dyad, we fit a generalized linear model (GLM) to obtain predictions $\hat{y}_j^{(i)}$ about the probability for each possible target collection $j \in \{1, 2, 3\}$ given the collection cycles i's covariates $\vec{x}^{(i)}$ and dyads fitted regression coefficients $\boldsymbol{\theta}$

$$\hat{y}_j^{(i)} = \text{Softmax}(\boldsymbol{\theta}^T \cdot \vec{x}^{(i)})_j, \tag{13}$$

where the softmax function is employed to transform the linear combination of covariates $\boldsymbol{\theta}^T \cdot \vec{x}^{(i)}$ into a probability distribution over the target choices $j \in \{1, 2, 3\}$

$$\text{Softmax}(\vec{x})_j = \frac{\exp(x_j)}{\sum_{k=1}^{3} \exp(x_k)}. \tag{14}$$

The covariates $\vec{x}^{(i)}$ for each collection cycle i consist out of three stacked components

$$\vec{x} = \begin{bmatrix} \vec{x}_d \\ \vec{x}_p \\ \vec{x}_{p'} \end{bmatrix}. \tag{15}$$

The first component $\vec{x}_d$ encodes the relevant distances to the targets as in section 1.1

$$\vec{x}_d = \begin{pmatrix} D_S(\vec{a}, \vec{b}, \vec{s}) \\ D_J(\vec{a}, \vec{b}, \vec{j}^A) \\ D_J(\vec{a}, \vec{b}, \vec{j}^B) \end{pmatrix}. \tag{16}$$

The second component $\vec{x}_p$ encodes the end result of the previous cycle as a one-hot-encoded vector. This encompasses: (1) which target was collected and if it was a single target, (2) whether there is an invite (as defined in section 1.3), and if there is an invite, then towards which of the two joint targets it is.

$$\vec{x}_p = \begin{cases} \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \end{pmatrix}^T & \text{if } \vec{s} \text{ previously collected and no invite,} \\ \begin{pmatrix} 0 & 1 & 0 & 0 & 0 \end{pmatrix}^T & \text{if } \vec{j}^A \text{ previously collected,} \\ \begin{pmatrix} 0 & 0 & 1 & 0 & 0 \end{pmatrix}^T & \text{if } \vec{j}^B \text{ previously collected,} \\ \begin{pmatrix} 0 & 0 & 0 & 1 & 0 \end{pmatrix}^T & \text{if } \vec{s} \text{ previously collected and invite towards } \vec{j}^A, \\ \begin{pmatrix} 0 & 0 & 0 & 0 & 1 \end{pmatrix}^T & \text{if } \vec{s} \text{ previously collected and invite towards } \vec{j}^B. \end{cases} \tag{17}$$

The third and last component encodes, also as a one-hot-encoded vector, which target type was collected two cycles ago.

$$\vec{x}_{p'} = \begin{cases} \begin{pmatrix} 1 & 0 & 0 \end{pmatrix}^T & \text{if } \vec{s} \text{ collected prior to the previously collected, i.e., two cycles ago,} \\ \begin{pmatrix} 0 & 1 & 0 \end{pmatrix}^T & \text{if } \vec{j}^A \text{ collected prior to the previously collected,} \\ \begin{pmatrix} 0 & 0 & 1 \end{pmatrix}^T & \text{if } \vec{j}^B \text{ collected prior to the previously collected.} \end{cases} \tag{18}$$

Analogous to the covariates $\vec{x}$ the coefficient matrix $\boldsymbol{\theta}$ consists out of three stacked components

$$\boldsymbol{\theta} = \begin{bmatrix} \boldsymbol{\theta}_d \\ \boldsymbol{\theta}_p \\ \boldsymbol{\theta}_{p'} \end{bmatrix}. \tag{19}$$

Where the first component $\boldsymbol{\theta}_d$ weights the relevant distance of the single target against those of the joint targets similar as the weighting w in section 1.1

$$\boldsymbol{\theta}_d = \begin{pmatrix} d_0 & 0 & 0 \\ 0 & d_1 & 0 \\ 0 & 0 & d_1 \end{pmatrix}. \tag{20}$$

The second component $\boldsymbol{\theta}_p$ does change either the relative probability of the previously collected target to the others or, if an invite is present, changes the relative probability of the joint target towards the invite is

$$\boldsymbol{\theta}_p = \begin{pmatrix} p_0 & 0 & 0 \\ 0 & p_1 & 0 \\ 0 & 0 & p_1 \\ 0 & p_2 & 0 \\ 0 & 0 & p_2 \end{pmatrix}. \tag{21}$$

The last component $\boldsymbol{\theta}_{p'}$ allows for a dependence further into the past. It modulates the probability of the next target towards the identity of the target prior to the previous target:

$$\boldsymbol{\theta}_{p'} = \begin{pmatrix} p_3 & 0 & 0 \\ 0 & p_3 & 0 \\ 0 & 0 & p_3 \end{pmatrix}. \tag{22}$$

Note that here in this component, the model does not differentiate between single and joint targets. To obtain the coefficient matrix $\boldsymbol{\theta}$, we use the quasi-Newton BFGS algorithm to minimize the negative log-likelihood function

$$\mathcal{L} = -\sum_{i=1}^{N} \sum_{j=1}^{3} y_j^{(i)} \ln(\hat{y}_j^{(i)}), \tag{23}$$

where $y_j^{(i)}$ is a binary encoding whether target collection j happened in collection cycle i or not and $\hat{y}_j^{(i)}$ is, as defined previously, the corresponding predicted probability of this target collection. The resulting prediction accuracy for unseen data (visualized in Fig. 4a) is evaluated for all collection cycles of the stable period (minutes 10–40) with k-fold cross-validation ($k = 5$). With this model, we can not only predict the dyad's target choice but we do also obtain an estimate of the dyad's uncertainty about the target choice at the start of the collection cycle (as utilized in Fig.5b and c). We estimate the dyads uncertainty with the uncertainty of our model which is quantified by the Shannon entropy H of the model's target prediction:

$$H(i) = -\sum_{j=1}^{3} \hat{y}_j^{(i)} \log_2(\hat{y}_j^{(i)}). \tag{24}$$

1.3 Trajectory classification

Trajectory classification is performed by a procedural algorithm that assigns a trajectory class (see Fig. 5a for visualizations) to each collection cycle based on a set of predefined conditions. The thresholds used in these conditions are set manually, and the trajectory classes are evaluated in a predetermined order. If a condition is met, the corresponding trajectory class is assigned to the cycle, and the algorithm proceeds to the next condition only if the current one is not met.

The first trajectory class evaluated is the "Invitation" class. A collection cycle is classified as "Invitation" if it follows a single target collection in the previous cycle, during which the non-collecting agent positioned itself in an "inviting" manner, that is, near a joint target before the collection of the previous single target collection is completed. This inviting placement can be either on the joint target (on-target-invite) or nearby the joint target (nearby-target-invite), with the latter defined as a distance less than half of that from the other agent. If the joint target is collected subsequently, the cycle is classified as "Invitation" otherwise, it is classified as "Failed invitation". If the nearby-target-invite condition is met, a cycle is only classified as "Failed invitation" if the overall FST is below 2/3, otherwise, it is more probable that the placement near the joint target is a coincidence. If neither the on-target-invite nor the nearby-target-invite condition is met, the algorithm proceeds to evaluate the "Strongly curved" trajectory class.

A collection cycle is classified as "Strongly curved" if the fraction of the excess trajectory length due to curvature exceeds 0.32. If this condition is not met, the agents' trajectories are relatively straight and the algorithm proceeds evaluating the condition of the "Different targets" trajectory class.
To determine whether the two agents aimed for different targets or the same, the algorithm fits straight lines to their trajectories, which serves to (1) smooth out the motor noise and (2) extend the aim of the agent beyond the final position reached. For each agent, the algorithm checks (1) whether the agent moved towards a target and (2) whether this target is in proximity to the corresponding line (i.e., less than 10.14 cm). If a target meets these conditions, it is considered a potential candidate for the agent's aim. If multiple candidate targets exist, the algorithm selects the one closest to the agent's final position after weighting the distances from the agents to the candidates according to the overall FST. This results in a prediction of which target each agent aimed for. If these predictions differ, the cycle is classified as "Different targets". Note that if both agents initially aimed at different joint targets, the collection cycle would not end, leading to curved trajectories, which is why different initial aims towards joint targets are subsumed under the "Strongly curved" trajectory class.
Finally, the algorithm checks whether one agent moves ahead towards the target at which both agents aim. If the difference in distance averaged over the collection cycle to the finally collected target exceeds 3.4 cm, the cycle is classified as "One ahead". If none of the above conditions hold true, the cycle is classified as "Concurrent" by exclusion principle.

1.4 Joint payoff

Table S3 | Variables used in section 1.4, and their descriptions.

Variable	Description
R	Joint payoff.
N	Number of target collections / collection cycles.
r	Payoff per target (7 cents).
T	Block duration (20 minutes).
t_i^{acq}	Duration of the acquisition period in cycle i.
$\langle\rangle$	Mean over all cycles.
t^{col}	Duration of the collection period (1 second).
e_1, e_2	Approximation errors.
l_i	Limiting trajectory length in cycle i.
s_i	Mean agent speed on the limiting trajectory in cycle i.
l_i^C	Additional trajectory length due to curvature in cycle i.
d_i	Limiting distance in cycle i.
d_i^T	Distance between the previous and the next target in cycle i.
d_i^R	Distance reduction due to advantageous placement in cycle i.

To elucidate the of role spatiotemporal factors in addition of those of classical discrete decision-making, we analyze in more detail the factors that determine the joint payoff of a dyad, using as a basis the observed trajectories of the dyads (Fig. 6). The joint payoff R equals the total number of target collections N times the payoff per target $r = 7$ cent

$$R = rN. \tag{25}$$

The total number of target collections N can be approximated by dividing the duration of each block $T = 20$ minutes by the mean collection cycle duration. The latter is the sum of the mean acquisition period duration $\langle t^{\mathrm{acq}} \rangle$ and the constant collection period duration $t^{\mathrm{col}} = 1$ second:

$$R = rN = r\frac{T}{\langle t^{\mathrm{acq}} \rangle + t^{\mathrm{col}}} + e_1. \tag{26}$$

The approximation error is denoted as e_1 and negligible ($r = 0.99, p < 10^{-6}$). The duration of an acquisition period t_i^{acq} equals the length of the so-called *limiting trajectory length* l_i divided by the corresponding speed on this trajectory s_i:

$$t_i^{\text{acq}} = \frac{l_i}{s_i}. \tag{27}$$

The length of the limiting trajectory l_i is either equal to the trajectory length of agent A or agent B. In case of a single target collection, it is that of the collecting agent, and in case of a joint target collection, it is that of the agent that entered last:

$$l_i = \begin{cases} l_i^A & \text{if } \vec{s}_i \text{ collected by agent } A, \\ l_i^B & \text{if } \vec{s}_i \text{ collected by agent } B, \\ l_i^A & \text{if } \vec{j}_i^X \text{ collected and agent } A \text{ entered last}, \\ l_i^B & \text{if } \vec{j}_i^X \text{ collected and agent } B \text{ entered last}. \end{cases} \tag{28}$$

By approximating the mean acquisition period duration $\langle t^{\text{acq}} \rangle$, we show that the joint payoff R is approximately proportional to the mean speed on and mean length of the limiting trajectory:

$$R = rN = r\frac{T}{\langle t^{\text{acq}} \rangle + t^{\text{col}}} + e_1 = r\frac{T}{\frac{\langle l \rangle}{\langle s \rangle} + t^{\text{col}}} + e_2. \tag{29}$$

The increased error, quantified by $e_2 > e_1$, originates from the inequality of the mean of a ratio to the ratio of means $\langle x/y \rangle \neq \langle x \rangle / \langle y \rangle$. Despite the increase, the error in predicting the joint payoff R remains negligible ($r = 0.99, p < 10^{-6}$). To decompose the joint payoff R further, we denote that the length of the limiting trajectory l is the sum of the excess length due to curvature l_i^C and the so-called *limiting distance* d_i.

$$l_i = l_i^C + d_i \tag{30}$$

The limiting distance d_i is the straight-line distance from the initial position of the limiting trajectory to the collected target. The limiting distance d_i results out of the position of the previously collected target and the amount of advantageous agent placement. Therefore, the limiting distance d_i is equal to the distance from the position of the previously collected target to the position of the subsequently collected target d_i^T minus the distance reduction due to advantageous placement d_i^R (see section 1.1)

$$l_i = l_i^C + d_i = l_i^C + d_i^T - d_i^R. \tag{31}$$

Substituting this into our joint payoff estimate, we observe its dependence on four variables:

$$R = rN = r\frac{T}{\langle t^{\text{acq}} \rangle + t^{\text{col}}} + e_1 = r\frac{T}{\frac{\langle l \rangle}{\langle s \rangle} + t^{\text{col}}} + e_2 = r\frac{T}{\frac{\langle l^C \rangle + \langle d^T \rangle - \langle d^R \rangle}{\langle s \rangle} + t^{\text{col}}} + e_2. \tag{32}$$

Here the error e_2 remains unchanged as the means of equal sample sizes are added. Two of the four variables (d^T and d^R) could in principle result out of classical discrete decision-making. The other two (l^C and s), however, result out of continuous, spatiotemporal interactions. Altogether, they shape the joint payoff in the cooperation–competition foraging game.

1.5 Individual payoffs

After analysing the joint payoff's dependencies, we proceed to decompose the individual payoff of a generic agent X and its counterpart, agent Y. We begin with two key relationships. We

Table S4 | Variables used in section 1.5, and their descriptions.

Variable	Description
X, Y	Generic agent names, could be either the blue or the orange target. If a variable is defined for agent X, the definition for agent Y is analogous.
R^X	Individual payoff of agent X.
R^X_Δ	Payoff difference of agent X to the other agent Y. Note that $R^X_\Delta = -R^Y_\Delta$.
n^X_s	Number of single target collections by agent X.
n^X_j	Number of joint target collections with a higher payoff share for agent X.
r_s	Payoff for a single target collection (7 cents).
$\hat{r}_j$	Higher payoff share of a joint target collection (5 cents).
$\check{r}_j$	Lower payoff share of a joint target collection (2 cents).
e	Error of predicting the payoff difference using only the single target difference. It is equal to the payoff difference due to differences in joint target type collections.

utilize (1) that the half of the joint payoff $R/2$ is exactly the within-dyad mean agent's individual payoffs (R^X and R^Y)

$$\frac{R^X + R^Y}{2} = \frac{R}{2} \tag{33}$$

and (2) that the difference from this intermediate payoff $R/2$ to the individual payoff R^X of an agent X equals the half of the inter-agent payoff difference $R^X_\Delta/2$ (defined later)

$$R^X = \frac{R}{2} + \frac{R^X_\Delta}{2}. \tag{34}$$

We observe that the individual payoff R^X depends on the payoff difference R^X_Δ and on the joint payoff R. From the previous section, we know on what variables the joint payoff R depends. Subsequently, we here analyze which variable is responsible for the inter-agent payoff difference

$$R^X_\Delta = R^X - R^Y. \tag{35}$$

The individual payoffs can be decomposed by the type of target:

$$R^X = n^X_s r_s + n^X_j \hat{r}_j + n^Y_j \check{r}_j \tag{36}$$

and

$$R^Y = n^Y_s r_s + n^Y_j \hat{r}_j + n^X_j \check{r}_j, \tag{37}$$

where n^X_s is the number of single target collections of agent X and $r_s = 7$ cent the payoff for a single target collection. n^X_j is the number of collected joint targets with a higher payoff share for agent X. This higher payoff share is $\hat{r}_j = 5$ cent. The lower joint target payoff share $\check{r}_j$ is 2 cents. The variables for agent Y are defined analogously. By substituting these calculations into our payoff difference formula we obtain:

$$R^X_\Delta = R^X - R^Y = (n^X_s - n^Y_s)r_s + (n^X_j - n^Y_j)(\hat{r}_j - \check{r}_j). \tag{38}$$

The right hand side of this equation has two components. The first, $(n^X_s - n^Y_s)r_s$, accounts for the payoff difference due to differences in single target collections. The second component, $(n^X_j - n^Y_j)(\hat{r}_j - \check{r}_j)$, accounts for the payoff difference due to difference in joint target collections. By leaving out the latter and introducing an error term e, we approximate the payoff difference

$$R^X_\Delta = R^X - R^Y = (n^X_s - n^Y_s)r_s + e \tag{39}$$

only with the difference in single targets. By demonstrating that the error e is negligible, we confirm that the inter-agent payoff difference is shaped predominantly by the difference in single targets (as visualized in Fig. 7b, $r = 0.99, p < 10^{-6}$). Thus, the individual payoff

R^X depends on the four variables known from previous section and on a fifth variable, the difference in single target collections between the agents.

1.6 Estimating the optimal strategy and cost of cooperation

Table S5 | Variables used in section 1.6, and their descriptions.

Variable	Description
Φ	Fraction of single targets (FST).
$\tilde{S}^X_\Delta$	Sensorimotor skill difference. It is the normalized single target difference S^X_Δ, considering only single target collections that are not due to defections like the "Different targets" and "Failed invitation" trajectory classes.
S^X_Δ	Normalized single target difference. It represents the difference in single target collections between agent X and agent Y, divided by the total number of single target collections. Note that $S^X_\Delta = -S^Y_\Delta$.
$\hat{R}^X(\Phi, \tilde{S}^X_\Delta)$	Estimated individual payoff of agent X given Φ and $\tilde{S}^X_\Delta$.
$\hat{R}(\Phi)$	Estimated joint payoff given Φ.
$\hat{R}^X_\Delta(\Phi, \tilde{S}^X_\Delta)$	Estimated inter-agent payoff difference for agent X given Φ and $\tilde{S}^X_\Delta$. Note that $\hat{R}^X_\Delta(\Phi, \tilde{S}^X_\Delta) = -\hat{R}^Y_\Delta(\Phi, \tilde{S}^Y_\Delta)$.
$\hat{\Phi}^X(\tilde{S}^X_\Delta)$	Estimated fraction of single targets of the optimal strategy of agent X given their skill difference $\tilde{S}^X_\Delta$.
$C(\Phi, \tilde{S}^X_\Delta)$	Estimated cost of cooperation for the higher-skilled agent X. It is called the cost of cooperation since the optimal strategy for a sufficiently higher-skilled agent is pure competition ($\Phi = 1$).

We show that participants are more cooperative than optimal and pay an associated cost of cooperation — the loss of income because of the choice of a non-optimal, overly cooperative, strategy. To estimate this cost of cooperation we estimate for each individual agent X the optimal strategy taking into account their estimated relative skill level. This optimal strategy is the strategy that maximizes the agent X's expected individual payoff. To predict this expected individual payoff, we model the individual payoff R^X of agent X as a function of FST Φ and skill difference $\tilde{S}^X_\Delta$ (defined later) analogous to the formula in section 1.5:

$$\hat{R}^X(\Phi, \tilde{S}^X_\Delta) = \frac{\hat{R}(\Phi)}{2} + \frac{\hat{R}^X_\Delta(\Phi, \tilde{S}^X_\Delta)}{2}, \tag{40}$$

where we use (1) an estimate of the inter-agent payoff difference $\hat{R}^X_\Delta(\Phi, \tilde{S}^X_\Delta)$ and (2) an estimate of the joint payoff $\hat{R}(\Phi)$. The latter follows out of an estimate of the number of target collections $\hat{N}(\Phi)$ which is analogous to that in section 1.4

$$\hat{R}(\Phi) = r\hat{N}(\Phi) = r\frac{T}{\frac{\hat{l}^C(\Phi)+\hat{d}^T(\Phi)-\hat{d}^R(\Phi)}{\hat{s}(\Phi)} + t^{\mathrm{col}}}. \tag{41}$$

The four dependencies, $\hat{l}^C(\Phi)$, $\hat{d}^T(\Phi)$, $\hat{d}^R(\Phi)$ and $\hat{s}(\Phi)$, of the joint payoff R are here estimated by fitting low-degree polynomials. This is a modeling choice to obtain later on the expected cost of cooperation for the average dyad. Alternatively, one could use here the values and assumptions from section 1.1 to obtain the cost of not pursuing the optimal dyad strategy. Having now defined the estimate for the joint payoff $\hat{R}(\Phi)$ we now continue with the estimate of the inter-agent payoff difference $\hat{R}^X_\Delta(\Phi, S^X_\Delta)$. Therefore, reusing the estimate from section 1.5:

$$R^X_\Delta = r_s(n^X_s - n^Y_s) + e \tag{42}$$

by multiplying $1 = \frac{N\Phi}{n_s^X + n_s^Y}$ we obtain

$$R_\Delta^X = r_s N\Phi \frac{n_s^X - n_s^Y}{n_s^X + n_s^Y} + e, \tag{43}$$

which motivates the definition of the normalized single target difference

$$S_\Delta^X = \frac{n_s^X - n_s^Y}{n_s^X + n_s^Y}. \tag{44}$$

This measure quantifies the efficiency of agent X in contrast to that of agent Y to collect single targets independently of FST. We only include single target collections when both agents moved straight to the single target ("concurrent" or "one ahead" to the same target), excluding defections ("Different targets" and "Failed invitation" trajectory classes from section 1.3) to obtain the competitive sensorimotor *skill difference* $\tilde{S}_\Delta^X$. This also makes this measure generalizable to different FST values. Now we define our new estimate of the inter-agent payoff difference as

$$\hat{R}_\Delta^X(\Phi, \tilde{S}_\Delta^X) = r\hat{N}(\Phi)\Phi\tilde{S}_\Delta^X, \tag{45}$$

where we utilize again our estimate of the number of target collections $\hat{N}(\Phi)$. Combining our estimate of the inter-agent payoff difference $\hat{R}_\Delta^X(\Phi, \tilde{S}_\Delta^X)$ with that of the joint payoff $\hat{R}(\Phi)$ we estimate agent X's individual payoff $\hat{R}^X(\Phi, \tilde{S}_\Delta^X)$ with high accuracy ($r = 0.88, p < 10^{-6}$) given only the FST Φ and the skill difference $\tilde{S}_\Delta^X$ (see Fig. 7c). Now we can estimate the FST value of agent X's optimal strategy

$$\hat{\Phi}^X(\tilde{S}_\Delta^X) = \arg\max_\Phi \hat{R}^X(\Phi, \tilde{S}_\Delta^X) \tag{46}$$

as well as the cost of cooperation of the higher-skilled agent X

$$C(\Phi, \tilde{S}_\Delta^X) = \max_\Phi(\hat{R}^X(\Phi, \tilde{S}_\Delta^X)) - \hat{R}^X(\Phi, \tilde{S}_\Delta^X) \tag{47}$$

for a given skill difference $\tilde{S}_\Delta^X$ (Fig. 7d).

Supplementary Figures

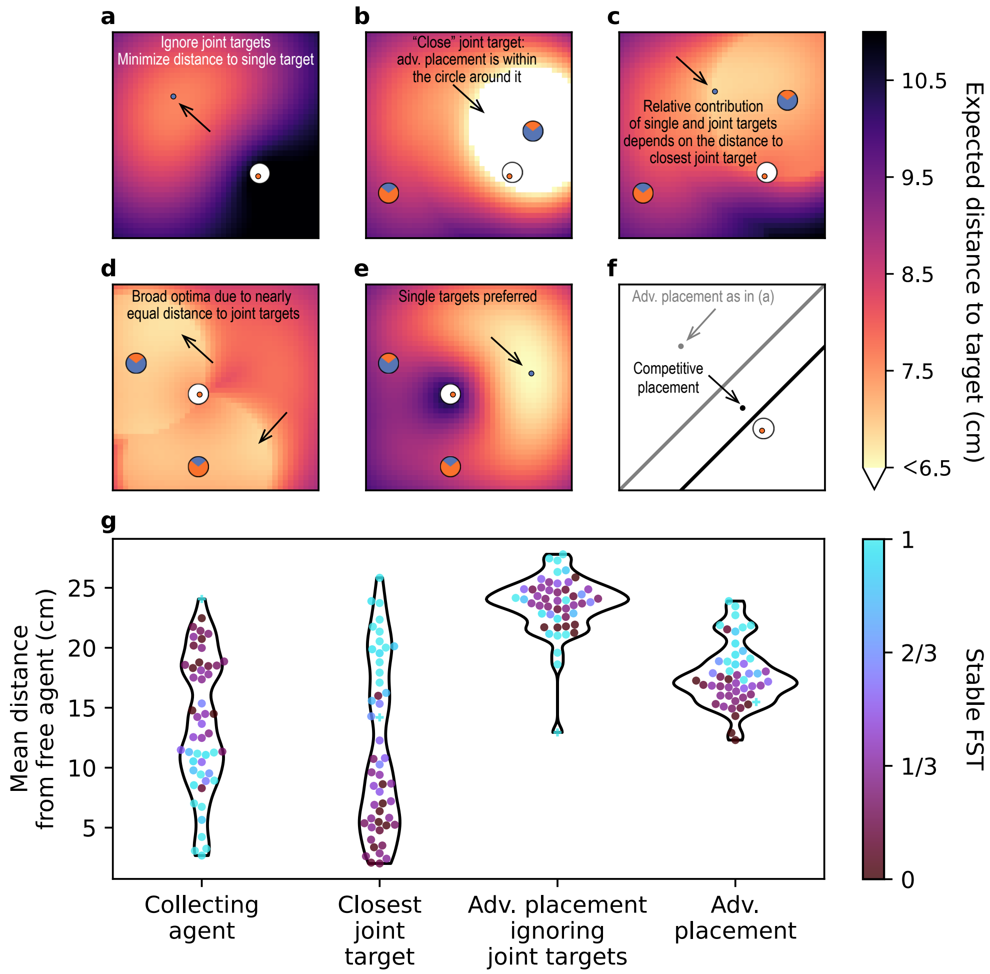

Fig. S1 | Advantageous and competitive placement, and actual placement during single target collections in each dyad. (**a-e**) Examples illustrating optimal (advantageous) placement of the non-collecting agent (blue dot, arrows) to minimize the expected distance to the next collected target (colormap). (**a**) Assuming no joint targets, the advantageous placement is on the other half of the game field, splitting the game field as much as possible between the two agents. Note that in (b-d), an equal weighting of target types is assumed (w=0.5). (**b**) When considering joint targets, the advantageous placement is within a circle centered on the closest joint target and extending to the currently collecting agent. (**c**) When the closest joint target is further away, the contribution of the single target becomes more apparent: the circle around the joint target is non-uniform and includes the optimum location splitting the game field between the two agents, similarly to (a). (**d**) In some spatial configurations, large parts of the game field are near optimal. (**e**) For a FST=1 dyad (w=0.99), the optimal placement ignores the joint targets (as in (a)). (**f**) Competitive placement example. The black dot represents a competitive placement, where the free, non-collecting agent optimizes the partitioning of the game field such that their probability of getting the next single target is maximized. For comparison, the gray dot represents the non-competitive advantageous placement as in (a). The gray line indicates the corresponding split of the game field into two parts, and the black line — the competitive split. (**g**) Actual mean distance from the free, non-collecting agent to four locations: (i) that of the collecting agent, (ii) that of the closest joint target, (iii) to the advantageous placement when ignoring the joint targets (w=0.99), and (iv) the advantageous placement (w=0.5). Instead of performing advantageous placement, agents in FST ≥ 0.9 dyads mainly perform competitive placement by placing themselves close to the collecting agent (as illustrated in (f)), and agents in $0.1 <$ FST < 0.9 dyads mostly place themselves next to the closest joint target. Note that FST ≤ 0.1 dyads are excluded from this analysis due to absent or very low number of single target collections.

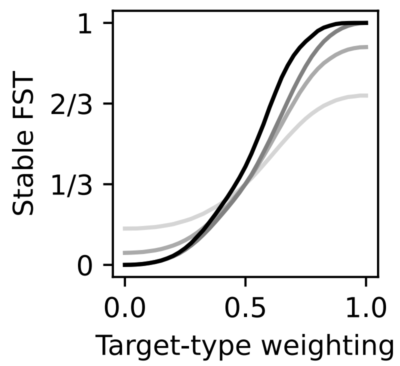

Fig. S2 | Cooperation–competition weighting of distances. To achieve different FSTs (fraction of single targets), the simulation requires different weightings of distances w reflecting specific dyads' preferences for selecting single targets vs. joint targets (eq. (3)). The black curve assumes an optimal "advantageous" placement of the non-collecting agent when the partner collects the single target, the dark gray curve assumes that the agents always move together and always select nearest (weighted) target, and the lighter gray lines assume non-optimal agents who move together but select a target randomly in 10% or 30% of trials (corresponding to the light gray lines in Fig. 3d).

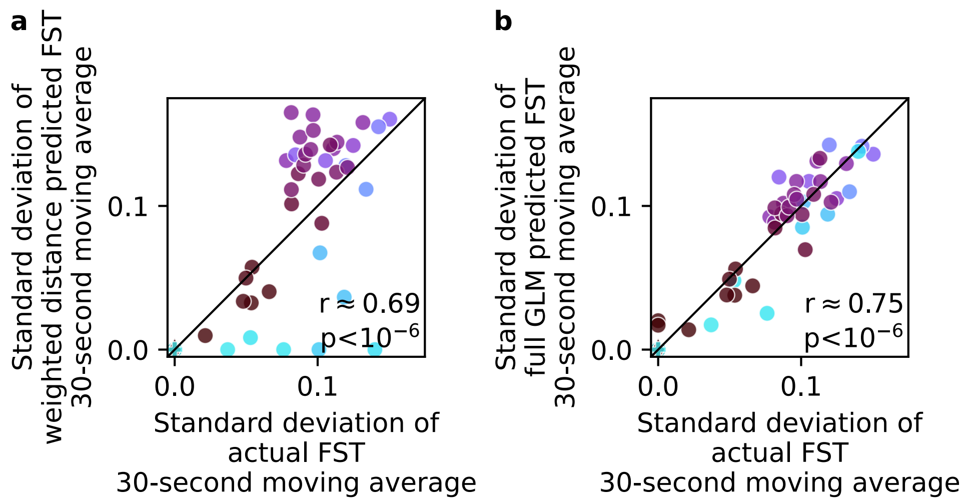

Fig. S3 | Variance of model predictions. (**a**) For high FST (light blue dots), the variance of predictions of the weighted distance model is too low whereas it is too high for intermediate FST (purple dots). (**b**) The full GLM is not only more accurate in its predictions (Fig. 4a), but also the variance of the predicted target type timecourse is better correlated with actual variance (Wilcoxon signed-rank test comparing the differences of standard deviations between 30-second moving average FST of the actual and each of the two model's predictions for the 40/58 dyads that exhibit FST fluctuations, $W = 258, p < 0.05, n = 40$, $\text{Mdn}_1 = -0.01\ [-0.03, 0.02]$, $\text{Mdn}_2 = 0.0008\ [-0.009, 0.02]$, $r_{rb} = 0.32, CI = [0.04, 0.62]$).

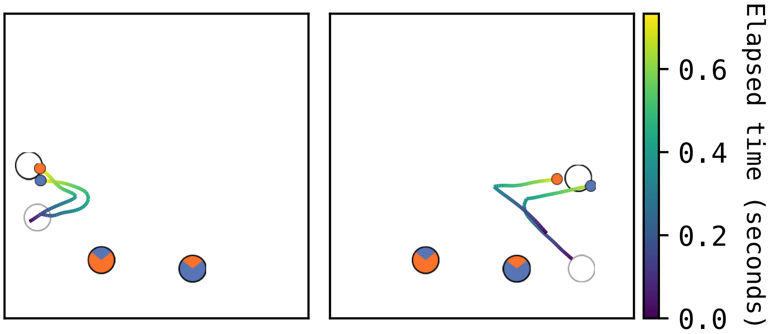

Fig. S4 | "Go-before-you-know" effect. Two example trajectories of a FST=1 dyad. Despite the certainty of target choice, the trajectories are strongly curved. Immediately after previous target collection, the agents move in the direction that minimizes the expected distance to the next target (the center of the game field), reflecting a preemptive strategy. Subsequently, once the newly appeared target is perceived, the trajectory is adjusted.

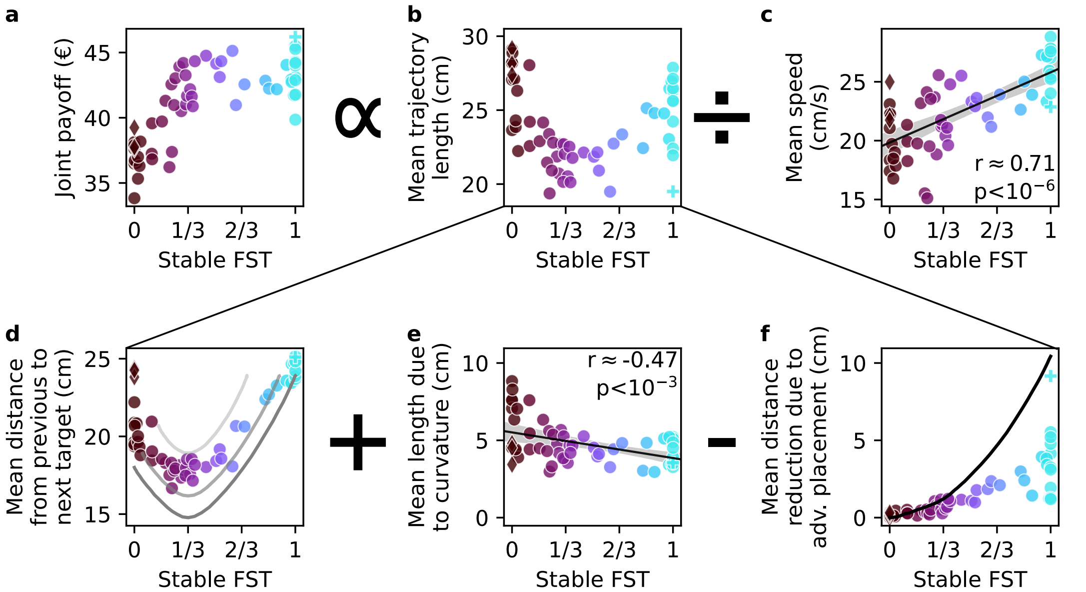

Fig. S5 | Spatiotemporal factors shape the payoff in a continuous action space. Note: panels (a), (b) and (c) are the same as in Fig. 6. (**a**) The joint payoff across both participants in a dyad is proportional to the mean acquisition duration. The mean acquisition duration is well-approximated by dividing the mean trajectory length (**b**) by the mean movement speed (**c**) on these trajectories. Note the increase in speed with higher fractions of single targets (FST). (**d-f**) We subdivide the mean trajectory length in (b) into three components. (**d**) Without advantageous placement, the trajectory length is at least the mean distance from the previous to the next target. The dark gray curve indicates the minimal mean distance attainable for each FST value, the lighter curves indicate the same path-minimizing strategy applied only in 90% and 70% of collection cycles. The highest distances are those of the three turn-taking dyads (diamond markers) who alternated between the two joint targets regardless of the distance. (**e**) Curved trajectories increase the mean trajectory length, especially at low FST because of frequent initial miscoordination between participants $(r(56) = -0.47, p < 10^{-3}, CI = [-0.65, -0.24]$. The exception to this pattern were the three turn-taking dyads, who eliminated miscoordination through their consistent strategy. (**f**) Advantageous placement reduces the trajectory length. To achieve this, the free agent must place itself strategically during a single target collection by the other. One dyad (plus marker) nearly reached the maximal attainable distance reduction (black curve, *cf.* Fig. 3d), using a *cooperative*, or at least a conflict-avoiding strategy, dividing the game field into a lower and upper half where each agent respectively collected the single targets.

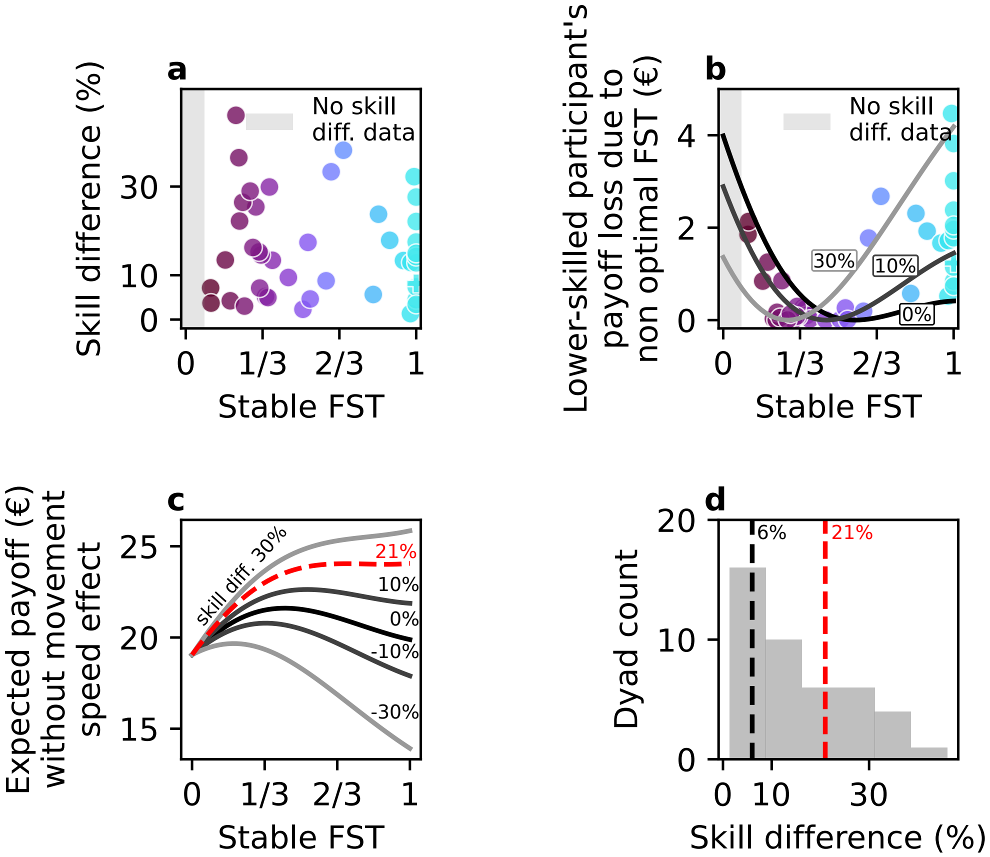

Fig. S6 | Skill differences and movement speed effect. (**a**) There is no observed relationship between skill difference in the dyad and the strategy they converged to. (**b**) Often the chosen strategy is also not optimal for the lower-skilled participant (*cf.* Fig. 7d). Note that each participant can increase FST on its own but decrease it only with the compliance of the other. (**c**) Expected payoff without the movement speed effect visualized as in Fig. 7c. The critical value of skill difference at which competition is optimal for the higher-skilled participant without taking into account the movement speed effects is at 21% (red dashed line). (**d**) Comparing the critical skill difference values expected with the movement speed effect (black dashed line) and without the speed effect (red dashed line). Due to the movement speed effect, competition is the best strategy for the majority of higher skilled participants.

Supplementary Movies

[YouTube playlist] | [Open Science Framework movies]

1. **Supplementary Movie S1:** YouTube or OSF
   Setup and game demonstration. Human Dyadic Interaction Platform setup and the game demonstration (60 s), followed by the replay of an example intermediate dyad.

2. **Supplementary Movie S2:** YouTube or OSF
   Cooperative example. Replay of a representative dyad from the cooperative group.

3. **Supplementary Movie S3:** YouTube or OSF
   Intermediate strategy example. Replay of a representative dyad from the intermediate group.

4. **Supplementary Movie S4:** YouTube or OSF
   Competition example. Replay of a representative dyad from the competitive group.

5. **Supplementary Movie S5:** YouTube or OSF
   Invitations examples. Collection cycles where one agent invites the other to a joint target.

6. **Supplementary Movie S6:** YouTube or OSF
   Cooperative turn-taking. Replay of one of three dyads who alternated between the two joint targets.

7. **Supplementary Movie S7:** YouTube or OSF
   Strongly curved trajectories. Examples of collection cycles featuring strongly curved trajectories, reflecting initial miscoordination and changes of mind.

8. **Supplementary Movie S8:** YouTube or OSF
   Competitive placement example. Replay of a dyad performing competitive advantageous placement.

9. **Supplementary Movie S9:** YouTube or OSF
   Cooperative placement for single targets. A special dyad that achieved nearly optimal advantageous placement by splitting the game field.

Instructions for participants

Welcome, and thank you for agreeing to participate in our experiment!

In this experiment you will play a game with another person through a transparent display.

Explanation of the game
In this game you need to collect targets to earn money. At the beginning the experimenter will assign a color to each player (blue and orange). On the screen you will see your own and other player's cursor (small circle colored respectively); you will be able to control your cursor with a computer mouse. The speed of the cursor is limited and by moving the mouse too fast you will have less control over it. To collect a target you need to place the cursor over the target and wait till the target completely disappears. At each time point of the game there will be 3 different targets available on the screen:

1. One white
2. One blue with a share of orange
3. One orange with a share of blue

You can collect white targets on your own — if you are first to select such a target, it becomes unavailable for the other player. To collect the colorful targets you need the other player to select the same target. Each collected target gives you a specific amount of money:

- White target gives 7 cents to the player who got it first
- Blue-orange target gives 5 cents to the blue player and 2 cents to the orange player
- Orange-blue target gives 5 cents to the orange player and 2 cents to the blue player

The amount of money you have collected will be displayed on the right side of the game field.

Structure of the experiment
The experiment will take ca. 2.5 hours in total. First, you will have a short game-tutorial to try out cursor control and target collection. Then, you will play 2 blocks, each 20 minutes, and between the blocks you will have an opportunity to take a break. From the start till the end of the experiment (including the breaks) we ask you to not communicate with the other player. In the end we will ask you to fill out a questionnaire.

Payment
You will receive the money you have collected in one of the blocks. In the end of the experiment, you will roll a dice to randomly select a block accordingly to which you will be paid:

- 1, 2, 3 - first block
- 3, 4, 5 - second block