



UvA-DARE (Digital Academic Repository)

Interpersonal trust: Asymptotic analysis of a stochastic coordination game with multi-agent learning

Meylahn, B.V.; den Boer, A.V.; Mandjes, M.

DOI

[10.1063/5.0205136](https://doi.org/10.1063/5.0205136)

Publication date

2024

Document Version

Final published version

Published in

Chaos

License

CC BY

[Link to publication](#)

Citation for published version (APA):

Meylahn, B. V., den Boer, A. V., & Mandjes, M. (2024). Interpersonal trust: Asymptotic analysis of a stochastic coordination game with multi-agent learning. *Chaos*, 34(6), Article 063119. <https://doi.org/10.1063/5.0205136>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Interpersonal trust: Asymptotic analysis of a stochastic coordination game with multi-agent learning

Cite as: Chaos 34, 063119 (2024); doi: 10.1063/5.0205136

Submitted: 24 February 2024 · Accepted: 16 May 2024 ·

Published Online: 7 June 2024



View Online



Export Citation



CrossMark

Benedikt V. Meylahn,^{1,a)} Arnoud V. den Boer,¹ and Michel Mandjes^{1,2}

AFFILIATIONS

¹Korteweg-de Vries Institute for Mathematics, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands

²Mathematical Institute, Leiden University, Niels Bohrweg 1, 2333 CA Leiden, The Netherlands

^{a)}Author to whom correspondence should be addressed: b.v.meylahn@uva.nl

ABSTRACT

We study the interpersonal trust of a population of agents, asking whether chance may decide if a population ends up with high trust or low trust. We model this by a discrete time, stochastic coordination game with pairwise interactions occurring at random in a finite population. Agents learn about the behavior of the population using a weighted average of what they have observed in past interactions. This learning rule, called an “exponential moving average,” has one parameter that determines the weight of the most recent observation and may, thus, be interpreted as the agent’s memory. We prove analytically that in the long run, the whole population always either trusts or doubts with the probability one. This remains true when the expectation of the dynamics would indicate otherwise. By simulation, we study the impact of the distribution of the payoff matrix and of the memory of the agents. We find that as the agent memory increases (i.e., the most recent observation weighs less), the actual dynamics increasingly resemble the expectation of the process. We conclude that it is possible that a population may converge upon high or low trust between its citizens simply by chance, though the game parameters (context of the society) may be quite telling.

© 2024 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0205136>

We are interested in a finite population playing a stochastic coordination game in random pairwise interactions, which we use to model an interaction requiring trust. Consider an agent’s point of view: “If others are abusing trust I prefer not to act trustingly while if others are honorable, then I would prefer to trust them.” Agents ideally coordinate on either trusting/acting honorably or on distrusting/abusing trust. The agents in our model are endowed with an adaptive learning rule. Models of learning in evolutionary game theory often assume that the agents’ opponent(s) are playing a fixed strategy and aim to learn this from experience. This makes little sense when all the agents are employing that learning rule (and, thus, updating their strategy). Our learning rule instead is an exponentially weighted moving average of the agents observations (also called simple exponential smoothing), which does not assume that the opponents’ strategy is fixed. Our model of a stochastic coordination game does away with fixed payoffs with noise and instead allows for a very broad category of random variables as payoffs with very mild conditions on their distributions. This is a new approach in the context of

agents who learn based on experience, though common in the study of the replicator equation. Another distinguishing element is that the stochastic game is dynamic as the payoffs are drawn at random anew at each time step. We prove convergence of belief and behavior in the long run to pure Nash equilibria (always: trust/act honorably or distrust/abuse trust). We conclude with simulations to explore the relationship between model parameters and relative probability of convergence to the trustful steady state and the rate of convergence. We see the surprising result that a shorter agent memory (i.e., more weight to the most recent observation) dampens the effect of the payoff distributions on the model outcomes.

I. INTRODUCTION

Trust is beneficial to societies’ functioning; yet, there is not a uniform amount of trust in different societies across the world.¹ Is there an inherent difference between the societies which exhibit

high trust and those that exhibit low trust? In Denmark, it is common practice to leave sleeping babies in their prams outside on the pavement, while its parents are shopping inside.² In contrast, one need to only look around and observe widespread electric fences and infrared alarm systems in a typical South African neighborhood to see that its citizens do not believe each other to be trustworthy.

What are the antecedents of low and high trust? Is there some insurmountable difference between the citizens of Denmark and those of South Africa, or is it possible that these differences are matters of chance? Furthermore, we ask what effect the structural properties of the interaction have on the probability of low or high trust emerging.

We apply the tool of evolutionary game theory with agent learning to the question of interpersonal trust. Instead of studying an N -person trust game³ as in Refs. 4–8, we make a simplification. In particular, we study a two-person coordination game played by two random agents drawn repeatedly from a population of N agents. This allows us to study what happens if both agents involved in the interaction learn from the interaction. This approach also illuminates the dynamics that occur for non-coordination games in which agents have positive externalities (utility derived from taking the same action as other players). We consider the problem of placing trust and acting in a trustworthy manner in society as a coordination game. We condense the matter of placing trust and honoring placed trust into the action *trust*. Similarly, not placing trust and abusing others' trust are condensed into the action *doubt*. Agents are happy to act trustingly while others are acting in a trustworthy manner, but prefer to act distrustfully when their neighbors are abusing trust; hence, the payoffs follow those of a coordination game. In the model, a population of agents interacts by random matching of two agents per round. These agents take the action that they believe to maximize their one-round expected utility. This expectation is based on their belief of “the probability that a random other in the population acts trustingly” and the randomly realized payoff in their stochastic coordination game. In this fictitious play-like⁹ model of learning, agents update their belief based on the exponential moving average rule.

We prove convergence of behavior and beliefs to the always-trust or always-doubt steady states in the long run. Furthermore, we highlight the impact that the exponential smoothing learning parameter has on the relative probability of convergence to the always-trust vs always-doubt steady state. Surprisingly, a greater learning rate (akin to a shorter memory of the agents) has a dampening effect on the impact of the game payoff parameter distributions. Conversely, based on our numerical simulations, we conjecture that a learning rate that approaches zero (an infinite memory) may create a phase transition between the always-trust and the always-doubt steady states depending on the distributions of the payoff variables.

Evolutionary game theory is a commonly used tool when it comes to understanding antecedents of social dynamics.^{10–16} In particular, models with agent learning have been studied extensively.^{17–26,67} Coordination games with slight payoff perturbations have been studied extensively as a tool to understand norm-formation.^{19,27–30} Our model is the first we know of which considers the case with fully stochastic payoffs rather than small random perturbations to a fixed payoff-matrix in the context of multi-agent learning.

We note the rich history of the fully stochastic payoff-matrix.^{31–42} In particular, there is a line of work on the statistics of the replicator (or the replicator–mutator) dynamics under the random draw of the game payoff-matrix.^{34,35} This is in contrast with Refs. 36–39, which study the replicator dynamics but in a changing environment (i.e., the payoff-matrix is drawn in each generation from some distribution). The replicator equation models a process of selection in which the reproductive fitness of a species (or strategy) is equal to its average payoff in the game. The object of interest in such studies is the number of (mixed) equilibria and their stability. Learning by imitation has also been studied in the context of a stochastic payoff-matrix.^{40–42} In such models, agents choose a neighbor and compare their most recent payoffs. If their neighbor has a greater payoff, the agent adopts their neighbors strategy with a given probability. In this branch of the literature, one looks for elements in the randomness or the population structure, which either promote or inhibit the emergence of trust.

The context of our model is experience based learning; i.e., the agents are learning about the population based on the strategies agents observe their opponents using and use this information to optimize their actions. Examples of this kind of learning with payoff matrices, which are either static or perturbed by small noise, can be found in Refs. 19, 27, and 43–46, respectively. The fact that agents converge to the always-trust or the always-doubt state complements (is in qualitative agreement with) the existing results of convergence to pure Nash equilibria⁴⁷ in the context of small perturbations and a decreasing learning rate in Refs. 17–19 and 48.

Our model illustrates that a *small* population may, by chance, end up in a low or high trust state. As the size of the population grows however, chance plays less of a role and the outcome is almost entirely decided by the distributions of their payoff parameters. The distributions of the payoff variables represent the context that a population is in. In situations where the payoff distributions promote doubting (low trust), it would behoove the individuals to have a shorter memory so that happy mistakes could lead to long-term trust.

Thus, because neither South Africa nor Denmark have “small” populations, our model leads us to believe that the low and high trust observed, respectively, may be largely due to the structural properties of interactions. By this, we mean the learning rate or the distributions of the payoffs in the coordination game. Our stylized mathematical model is of a high level of abstraction and aims to capture the essence of the underlying dynamics. Evidently, in the (inherently more complex) real dynamics, there are various other differences between South Africa and Denmark.

In Sec. II, we formulate the model formally, which we follow in Sec. III with an asymptotic analysis. This culminates in the main theorem regarding convergence of belief and behavior of the agents. In Sec. IV, we provide an illustration of the model with only two agents, which highlights the perhaps surprising nature of our main results. Having identified the asymptotic behavior of the dynamics, we introduce and discuss the results of a simulation study in Sec. V. The simulation serves to elucidate the interdependence of the parameters on the chance of convergence to high or low trust behavior. We conclude in Sec. VI with a discussion of the results and possible future work.

II. MODEL

We consider a population of $N \in \mathbb{N}$ ($N \geq 2$) agents who engage in a game played repeatedly in discrete rounds indexed by $t \in \mathbb{N}$. \mathbb{N} represents the natural numbers, $1, 2, 3, \dots$ in this paper. At the start of each round $t \in \mathbb{N}$, a pair of agents $(I(t), J(t))$ is chosen uniformly at random from the set of tuples $\{(i, j) \in \{1, \dots, N\} : i \neq j\}$. This may either be interpreted as if the population is fully mixed and there is no population structure or, equivalently, as if the structure imposed on the population is the fully connected network. The chosen pair of agents play a 2×2 coordination game. In this game, $I(t)$ takes the role of agent $k = 1$ (the row agent) and $J(t)$ takes the role of agent $k = -1$ (the column agent). We define $g(k, t)$,

$$g(k, t) := \begin{cases} I(t) & \text{if } k = 1 \text{ and} \\ J(t) & \text{if } k = -1 \end{cases} \quad \text{for } t \in \mathbb{N}.$$

This allows us to retrieve an agent's index within the population given their role in the game. The action of agent $k \in \{1, -1\}$ at time $t \in \mathbb{N}$ is denoted $A_k(t) \in \{T, D\}$. The action T denotes trusting and D denotes the action of doubting. We define the payoff bimatrix of the game as $\Pi(t)$ at time $t \in \mathbb{N}$,

$$\Pi(t) = \begin{matrix} & \begin{matrix} T & D \end{matrix} \\ \begin{matrix} T \\ D \end{matrix} & \begin{bmatrix} (U_1(t), U_{-1}(t)) & (V_1(t), W_{-1}(t)) \\ (W_1(t), V_{-1}(t)) & (Y_1(t), Y_{-1}(t)) \end{bmatrix} \end{matrix}. \quad (1)$$

The first value of the matrix entry $\Pi_{l,m}(t)$ is the reward obtained by the row agent when playing the action in row l against a column agent who plays the action in column m during round $t \in \mathbb{N}$. Conversely, the second value of the matrix entry $\Pi_{l,m}(t)$ is the reward obtained by the column agent when playing the action in column m against a row agent who plays the action in row l during round $t \in \mathbb{N}$. We define $\Pi_k(t)$ as the matrix containing only the payoffs to agent $k \in \{-1, 1\}$ at time $t \in \mathbb{N}$.

We assume that the payoffs in (1) $\{U_k(t) : k \in \{1, -1\}, t \in \mathbb{N}\}$ are all independent and identically distributed continuous random variables. The same is true for $\{V_k(t) : k \in \{1, -1\}, t \in \mathbb{N}\}$, $\{W_k(t) : k \in \{1, -1\}, t \in \mathbb{N}\}$, and $\{Y_k(t) : k \in \{1, -1\}, t \in \mathbb{N}\}$. To ensure that the game has the structure of a coordination game, we assume that

$$U_k(t) > V_k(t) \quad \text{and} \quad U_k(t) > W_k(t) \quad \text{w.p. } 1 \quad (2)$$

and

$$Y_k(t) > V_k(t) \quad \text{and} \quad Y_k(t) > W_k(t) \quad \text{w.p. } 1 \quad (3)$$

for $k \in \{1, -1\}$ and all $t \in \mathbb{N}$.

Remark 1. The restrictions (2) and (3) are stronger than what we need for our results. Our analysis holds whenever

$$Y_k(t) - V_k(t) > 0,$$

$$U_k(t) - W_k(t) \geq 0,$$

and

$$U_k(t) - W_k(t) + Y_k(t) - V_k(t) > 0$$

for all $t \in \mathbb{N}, k \in \{1, -1\}$. We add the assumptions (2) and (3) to show the relevance of the analysis to a coordination game.

In our setting, the agents model the population behavior with a belief on the probability that a randomly chosen individual that

they play against would trust. Let $x_i(t)$ denote the belief of agent $i \in \{1, \dots, N\}$ at the beginning of round $t \in \mathbb{N}$ on the likelihood that they will encounter an agent playing trust. The vector holding the beliefs of the agents in the population at time t is $\mathbf{x}(t) = (x_1(t), \dots, x_N(t))^T \in [0, 1]^N$.

At the start of each round, each agent $k \in \{1, -1\}$ only observes their own payoffs $U_k(t), Y_k(t), V_k(t), W_k(t)$ but not the payoffs of their opponent $U_{-k}(t), Y_{-k}(t), V_{-k}(t), W_{-k}(t)$. Furthermore, the agents are not aware that their opponent's payoffs follow the structure of a coordination game. The agents model their opponent's behavior entirely by their belief on the likelihood that their opponent take the trust action. We define $u_k^T(t)$ and $u_k^D(t)$ as the expected utility for agent $k \in \{1, -1\}$ playing T and D , respectively, during round $t \in \mathbb{N}$ based on agent k 's belief,

$$u_k^T(t) := x_{g(k,t)}(t)U_k(t) + (1 - x_{g(k,t)}(t))V_k(t), \quad (4)$$

$$u_k^D(t) := x_{g(k,t)}(t)W_k(t) + (1 - x_{g(k,t)}(t))Y_k(t). \quad (5)$$

Myopic decision making is a common assumption.⁴⁹⁻⁵⁴ Furthermore, there is experimental work that suggests that humans indeed act at least semi-myopically.^{55,56} As such, we assume the agents take actions myopically,

Assumption 1 (Myopic rationality). We assume agents to be myopically rational, taking the action, which maximizes the one-round expected utility, with ties favoring trust,

$$A_k(t) := \begin{cases} T & \text{if } u_k^T(t) \geq u_k^D(t), \\ D & \text{else} \end{cases}$$

for $t \in \mathbb{N}$ and $k \in \{1, -1\}$.

By letting ties favor trust, we obtain the weak inequality. We do this to facilitate clean analysis. It should be noted, however, that event in which equality holds has probability zero because the payoffs are continuous random variables. We will often be interested in the case when both agents take the same action. As such, we define $A(t)$ for $t \in \mathbb{N}$ without subscript as

$$A(t) := \begin{cases} T & \text{if } A_{-1}(t) = A_1(t) = T, \\ D & \text{if } A_{-1}(t) = A_1(t) = D, \\ 0 & \text{else.} \end{cases}$$

When modeling agents that learn in the context of game theory, there is a variety of learning mechanisms that may be implemented. For example, the agents could implement a basic form of reinforcement learning, which would be an adjustment to our Assumption 1, combining belief and action on belief into one process. Alternatively, agents could make repeated application of Bayes' rule given the interaction they observe or use a stochastic approximation algorithm⁵⁷ to update their belief on the probability that a random opponent will trust. The agents in our model use the exponential moving average⁵⁸ of their experiences. This is a straightforward, but powerful rule of thumb applied in signal processing. We introduce $\bar{\alpha} := 1 - \alpha$ for improved legibility of future equations.

Assumption 2 (Belief updating). The agents that are selected to play a game in round t update their belief based on the outcome of that game using exponential smoothing with learning

rate $\alpha \in (0, 1)$,

$$x_{g(k,t)}(t + 1) := \bar{\alpha}x_{g(k,t)}(t) + \alpha \mathbb{1}\{A_{-k}(t) = T\}$$

for $k \in \{1, -1\}$ and $t \in \mathbb{N}$. All other agents $i \in \{1, \dots, N\} \setminus \{I(t), J(t)\}$ retain their most recent belief,

$$x_i(t + 1) := x_i(t), \quad t \in \mathbb{N}.$$

We make this assumption mainly in order to facilitate clean analysis. Furthermore, we believe that this belief updating rule is simple enough for it not to be overly unrealistic to assume that people might learn in a similar fashion.

We note the work of Sato and collaborators,^{43–45} which also features updates according to the exponential moving average. An important difference between that work and ours is that we consider the discrete and random dynamics while they make use of a separation of time scales, which forces actual dynamics to resemble the expectation. This last element is something, we shall see, emphatically not present in our model.

The learning rate α is the weight of an agent’s most recent observation and may, thus, be interpreted as the agent’s memory. A greater weight to the most recent observation means that the earlier observations weigh less and, thus, are more quickly forgotten.

To an outsider who has not observed $\Pi_k(t)$ for $k \in \{1, -1\}$ during round $t \in \mathbb{N}$, the action taken looks random. Specifically, if the outsider is privy to the distributions of U_k, Y_k, V_k , and W_k , as well as the belief $x_k(t)$, then the probability that agent $k \in \{1, -1\}$ plays trust in round $t \in \mathbb{N}$ is defined as $p_k(t)$,

$$p_k(t) := \mathbb{P}(u_k^T(t) \geq u_k^D(t)). \tag{6}$$

To clarify, for the agent $k \in \{-1, 1\}$ who knows the values in $\Pi_k(t)$ and their own belief $x_k(t)$, the truth of the inequality $u_k^T(t) \geq u_k^D(t)$ is not random. For an outsider who is aware of the belief $x_k(t)$ and the distributions but not the realizations of the payoffs, the agent’s actions seem random and follow the probability defined in (6). The

behavior defined in (6) may be restated for $k \in \{1, -1\}$ and $t \in \mathbb{N}$,

$$p_k(t) = \mathbb{P}(Z_k(t) \leq x_k(t)),$$

for which we define $Z_k(t)$ as the random variable,

$$Z_k(t) := \frac{Y_k(t) - V_k(t)}{U_k(t) + Y_k(t) - V_k(t) - W_k(t)} \tag{7}$$

for $k \in \{-1, 1\}, t \in \mathbb{N}$. We acknowledge the freedom we have in defining a cdf for $Z_k(t)$ by allowing correlation between its constituent random variables. Using this freedom, we continue the discussion and analysis using $Z_k(t)$ and defining its cdf,

$$F(x) := \mathbb{P}(Z_k(t) \leq x), \quad x \in \mathbb{R}, \tag{8}$$

which then also defines agent behavior $p_k(t) = F(x_k(t))$ for $k \in \{1, -1\}$ and $t \in \mathbb{N}$. We make the following assumption on F , which is essentially of a technical nature.

Assumption 3 (Assumptions on F). We consider only $F(\cdot)$ for which it is true that F is a cumulative distribution function (non-decreasing and right-continuous) with $F(0) = 0, F(1) = 1$, and $F(x) \in (0, 1)$ for all $x \in (0, 1)$.

This assumption implies that if an agent believes that their opponent will trust (or doubt) with probability 1, that they too will trust (or doubt) with probability 1. Furthermore, that $F(\cdot)$ is a cdf and, therefore, non-decreasing, means that an agent with a higher belief is at least as likely to trust than an agent with a lower belief. Because $U_k(t), Y_k(t), V_k(t)$, and $W_k(t)$ are, respectively, iid between rounds $t \in \mathbb{N}$ as well as players $k \in \{-1, 1\}$, the cdf $F(\cdot)$ does not change from round to round or from player to player.

The model we have described is highlighted by the flow chart in Fig. 1. This shows the process that the two randomly chosen agents follow in one given round. We are interested in the evolution of the belief vector $\mathbf{x}(t)$ for $t \in \mathbb{N}$. Specifically, we ask whether there is convergence of beliefs to an equilibrium state.




FIG. 1. A flow chart of the model. This shows the process which the pair of randomly selected agents go through for one round. In each round, two agents are drawn from the population at random to play the stochastic coordination game once against each other. Each of these agents observes an independent draw of the random payoff matrix and compares this with their belief. If their expected reward for doubting is greater than that for trust, then they doubt and otherwise they trust. Then, each of these agents adjusts its belief based on the action taken by its respective opponent and returns to the pool of possible agents to be selected in the next round.

III. ASYMPTOTIC ANALYSIS

In this section, we analyze the long-term dynamics of the process in the limit as time $t \rightarrow \infty$. We first show that it is possible for the N agents to absorb (never exit an ϵ -ball around) at $\mathbf{x} = \mathbf{0}$ and $\mathbf{x} = \mathbf{1}$ where agent behavior converges to always-doubt and always-trust, respectively. Subsequently, we show that the process may end up in these corners from the interior of the state space. Finally, we combine these two sub-results in a Borel–Cantelli argument, which proves that the process will converge to one of these corners in the long run with probability one. On an intuitive level, this is explainable by the fact that in a coordination game, the agents will “try” to coordinate their behavior. Thus, under learning and myopic rationality, it is rational for a population to always trust or always doubt as this guarantees 100 percent coordination.

We state the main theorem, and subsequently, we will state and prove the lemmas required in the proof of this main theorem. We end this section by providing the proof of Theorem 1.

In the statement of the main theorem, we use Lagrange notation to define $F^{(n)}(x)$ as the n th derivative of F at x .

Theorem 1 (Absorption in a corner is guaranteed). *If $\alpha > \epsilon$ and $F^{(n_1)}(0), F^{(n_2)}(1) \in [0, \infty)$ for some $n_1, n_2 < \infty$, then*

$$\mathbb{P}(\exists t_0 : \{\mathbf{x}(t) \in [0, \epsilon]^N, \forall t > t_0\} \cup \{\mathbf{x}(t) \in [1 - \epsilon, 1]^N, \forall t > t_0\}) = 1, \quad (9)$$

which may be written as

$$\mathbb{P}(\exists t_0 : \mathbf{x}(t) \in [0, \epsilon]^N \cup [1 - \epsilon, 1]^N, \forall t > t_0) = 1.$$

The conditions on F in this theorem are mild. All that we require of F is that there exist finite integers n_1 and n_2 such that the n_1 -th derivative of F at $x = 0$ exists and is finite and that the n_2 th derivative of F at $x = 1$ exists and is finite.

A. Absorption is possible

The first lemma we will need pertains to the possibility of a population converging in an ϵ -ball around zero:

Lemma 1 (Absorption at zero for N agents). *Let $\alpha \in (0, 1)$, $\epsilon \in (0, \alpha)$, and suppose $\mathbf{x}(t_0) \in [0, \epsilon]^N$ for some $t_0 \in \mathbb{N}$, then*

$$\mathbb{P}(\mathbf{x}(t) \in [0, \epsilon]^N, \forall t \geq t_0 \mid \mathbf{x}(t_0) \in [0, \epsilon]^N) > 0. \quad (10)$$

In most of the probabilities that we write, there is a condition on the state of the system at t_0 . In order to fit equations into the available space as well as for general legibility, we define the following notation:

$$\mathbb{P}_{[a,b]^N}(\cdot) := \mathbb{P}(\cdot \mid \mathbf{x}(t_0) \in [a, b]^N). \quad (11)$$

Proof. In the new notation, we will show that

$$\mathbb{P}_{[0,\epsilon]^N}(\mathbf{x}(t) \in [0, \epsilon]^N, \forall t \geq t_0) > 0. \quad (12)$$

We prove this by induction. As a base case for our induction, we prove that absorption is possible for $N = 2$,

$$q := \mathbb{P}_{[0,\epsilon]^2}(\mathbf{x}(t) \in [0, \epsilon]^2, \forall t \geq t_0) > 0. \quad (13)$$

When $\mathbf{x}(t) \in [0, \epsilon]^2$, then even one agent playing T in some round $t_1 \in \mathbb{N}$, $t_1 \geq t_0$ implies that $\mathbf{x}(t_1 + 1) \notin (0, \epsilon)^2$ because

$x\bar{\alpha} + \alpha > \alpha > \epsilon$ for all $x \in (0, 1)$ and $\alpha \in (0, 1)$ and any $\epsilon \in (0, \alpha)$. Thus, the probability of remaining in $[0, \epsilon]^2$ for all $t = t_0, t_0 + 1, \dots, t_0 + n$ is the probability of both agents playing doubt in all rounds from t_0 until (and including) round $t_0 + n$,

$$\begin{aligned} &\mathbb{P}_{[0,\epsilon]^2}(\mathbf{x}(t) \in (0, \epsilon)^2, t = t_0, \dots, t_0 + n) \\ &= \mathbb{P}_{[0,\epsilon]^2}(A(t) = D, t = t_0, \dots, t_0 + n). \end{aligned} \quad (14)$$

Note that we are interested in the limit of the above expression as $n \rightarrow \infty$ because the agents are required to play D for all future rounds. The probability that agent 1 plays D in round t is given by $1 - F(x_1(t))$. For legibility, we define $\bar{F}(\cdot) := 1 - F(\cdot)$. Because the agents' actions during any round are independent, we have the above probability restated as

$$\begin{aligned} &\mathbb{P}_{[0,\epsilon]^2}(A(t) = D, t = t_0, \dots, t_0 + n) \\ &= \prod_{m=0}^n \bar{F}(\bar{\alpha}^m x_1(t_0)) \bar{F}(\bar{\alpha}^m x_2(t_0)). \end{aligned} \quad (15)$$

In order to simplify an expression, we define

$$D_n := \prod_{m=0}^n \bar{F}(\bar{\alpha}^m x_1(t_0)) \bar{F}(\bar{\alpha}^m x_2(t_0)). \quad (16)$$

Let $z_0 := \max\{x_1(t_0), x_2(t_0)\}$ and subsequently $z_m := \bar{\alpha}^m z_0 \forall m \in \mathbb{N}$. We bound (15) by

$$D_n \geq \prod_{m=0}^n \bar{F}(z_m)^2$$

because $z_m \geq x_i(t_0)\bar{\alpha}^m$ for all $m \in \mathbb{N}$ and because $F(x) \geq F(y)$ as long as $x \geq y$ by virtue of being a distribution function. We take the logarithm of both sides as well as the limit as $n \rightarrow \infty$,

$$\lim_{n \rightarrow \infty} \log(D_n) \geq \lim_{n \rightarrow \infty} \log\left(\prod_{m=0}^n (\bar{F}(z_m))^2\right),$$

and change the logarithm of a product on the right to a sum of logarithms,

$$\lim_{n \rightarrow \infty} \log(D_n) \geq \lim_{n \rightarrow \infty} 2 \sum_{m=0}^n \log(\bar{F}(z_m)). \quad (17)$$

We intend to bound the right-hand side of (17). To do this, observe first that $e^{-x} \geq 1 - x$. This may be checked by evaluating the tangent of e^{-x} at $x = 0$ and noting that e^{-x} is convex and so remains above this tangent line. We rearrange to obtain

$$1 - e^{-x} \leq x, \quad \forall x \in \mathbb{R}.$$

We substitute $x = \log(y)$ to get

$$1 - e^{-\log(y)} \leq \log(y),$$

which simplifies to

$$1 - \frac{1}{y} \leq \log(y), \quad \forall y \in \mathbb{R}_{>0}. \quad (18)$$

We use the above inequality with $y = \bar{F}(z_m)$ to bound each term in the sum from below by

$$\log(\bar{F}(z_m)) \geq 1 - \frac{1}{\bar{F}(z_m)} = \frac{F(z_m)}{F(z_m) - 1}.$$

Implanting this into the sum in (17) gives

$$\lim_{n \rightarrow \infty} \log(D_n) \geq \lim_{n \rightarrow \infty} 2 \sum_{m=0}^n \frac{F(z_m)}{F(z_m) - 1}. \quad (19)$$

In order to invoke Abel's convergence test (see, for example, Ref. 59), we note that the sum in (19) is the product of the sequences $\{F(z_m)\}$ and $\{1/(F(z_m) - 1)\}$. The second of these is bounded from above by -1 and is monotone increasing in m . This is because $F(z_m)$ is monotone decreasing in m and bounded by 0. Therefore, if $\sum_{m=0}^{\infty} F(z_m)$ converges, then the right hand side of (19) also converges.

To prove convergence of $\sum_{m=0}^{\infty} F(z_m)$, we use the ratio test. We will show that

$$\lim_{m \rightarrow \infty} \left| \frac{F(z_{m+1})}{F(z_m)} \right| = \lim_{m \rightarrow \infty} \left| \frac{F(\bar{\alpha}^{m+1} z_0)}{F(\bar{\alpha}^m z_0)} \right| < 1. \quad (20)$$

The inequality is a result of an (possibly repeated) application of L'Hôpital's rule. First note that

$$\lim_{m \rightarrow \infty} F(z_m) = \lim_{m \rightarrow \infty} F(z_{m+1}) = 0,$$

and therefore, the limit of the quotient in (20) yields an indeterminate form 0/0. We apply L'Hôpital's rule,

$$\lim_{m \rightarrow \infty} \left| \frac{F(z_{m+1})}{F(z_m)} \right| = \bar{\alpha} \frac{\lim_{m \rightarrow \infty} F'(\bar{\alpha}^{m+1} z_0)}{\lim_{m \rightarrow \infty} F'(\bar{\alpha}^m z_0)}.$$

- Either at this point, we have that $\sum_k F(z_m)$ converges if $F'(0) = c \in (0, \infty)$ or
- we apply L'Hôpital's rule n times until $F^{(n)}(0) = c \in (0, \infty)$ at which point the limit is $\bar{\alpha}^n < 1$.

Thus, all conditions for Abel's convergence test are satisfied, and we have shown convergence of the right hand side of (19). In particular, this will converge to a negative number because all the terms are negative as a result of the denominator. By taking the exponential of both sides of (19), we get the probability of both agents playing doubt indefinitely on the left. On the right hand side, we get a positive number. Considering that the probability of convergence is equal to the probability of both agents playing doubt indefinitely, we have shown that this positive,

$$\mathbb{P}_{[0, \epsilon]^2}(\mathbf{x}(t) \in [0, \epsilon]^2, \forall t \geq t_0) = \lim_{n \rightarrow \infty} D_n > 0.$$

This proves our base case: $q > 0$. Our induction hypothesis, thus, states that N players can converge. As such, we denote by q_N the

probability of N players converging in an ϵ -ball at zero,

$$q_N := \mathbb{P}_{[0, \epsilon]^N}(\mathbf{x}(t) \in [0, \epsilon]^N, \forall t \geq t_0) > 0. \quad (21)$$

Now, we intend to show that

$$\mathbb{P}_{[0, \epsilon]^{N+1}}(\mathbf{x}(t) \in [0, \epsilon]^{N+1}, \forall t \geq t_0) > 0. \quad (22)$$

We identify (arbitrarily) the first agent and collect all the rounds in which agent $i = 1$ is selected to play in M ,

$$M := \{t \geq t_0 : 1 \in \{I(t), J(t)\}\}.$$

We note that as before, if at least one agent plays trust in any round t , then $\mathbf{x}(t + 1) \notin [0, \epsilon]^{N+1}$. Thus, we have the relationship

$$\begin{aligned} \mathbb{P}_{[0, \epsilon]^{N+1}}(A(t) = D, \forall t \geq t_0) \\ = \mathbb{P}_{[0, \epsilon]^{N+1}}(\mathbf{x}(t) \in [0, \epsilon]^{N+1}, \forall t \geq t_0). \end{aligned} \quad (23)$$

We may write the probability of all agents playing D persistently as a product of conditional probabilities as follows:

$$\begin{aligned} \mathbb{P}_{[0, \epsilon]^{N+1}}(A(t) = D, \forall t \geq t_0) &= \prod_{j=0}^{\infty} \mathbb{P}_{[0, \epsilon]^{N+1}}(A(t_0 + j) \\ &= D \mid A_k(\tau) = D, \forall t_0 \leq \tau < t_0 + j). \end{aligned} \quad (24)$$

Because we are dealing with multiplication on the right hand side, the ordering of terms does not matter. We are, thus, free to collect all the terms involving agent 1 in one product (a) and the remaining terms in a different product (b),

$$\begin{aligned} \mathbb{P}_{[0, \epsilon]^{N+1}}(A(t) = D, \forall t \geq t_0) \\ = \underbrace{\prod_{t \in M} \mathbb{P}_{[0, \epsilon]^{N+1}}(A(t) = D \mid A(\tau) = D, \forall t_0 \leq \tau < t)}_{=: a} \\ \times \underbrace{\prod_{t \notin M} \mathbb{P}_{[0, \epsilon]^{N+1}}(A(t) = D \mid A(\tau) = D, \forall t_0 \leq \tau < t)}_{=: b}. \end{aligned} \quad (25)$$

For the rounds $t \in M$, we introduce the notation $t_{i,l}$ to be the time of the l th round in which agent 1 is chosen to play against agent i . Note that because the dynamics never end, agent one will be chosen to play with each other agent for an infinite number of rounds. Thus, we can split a further and in terms of agent beliefs,

$$a = \prod_{i=2}^{N+1} \prod_{l=1}^{\infty} \bar{F}(x_1(t_{i,l})) \bar{F}(x_i(t_{i,l})).$$

We can bound the beliefs of the agents involved because we know that at least they have played against each other l times and always played doubt (by the conditioning). Therefore, we set $\tilde{x}_l := \epsilon \bar{\alpha}^{l-1}$ for $l = 1, 2, \dots$ and note that $x_j(t_{i,l}) \leq \tilde{x}_l$ for both $j = 1$ and i as well as all $l = 1, 2, \dots$

We can bound the probability in the product by

$$a \geq \prod_{i=2}^{N+1} \prod_{l=1}^{\infty} \bar{F}(\tilde{x}_i)^2 = \prod_{l=1}^{\infty} (\bar{F}(\tilde{x}_i))^{2N}.$$

We take the logarithm on both sides,

$$\log(a) \geq 2N \sum_{l=1}^{\infty} \log(\bar{F}(\tilde{x}_i)).$$

From here, we can repeat the steps followed in the proof of our base case from (17) to conclude that the sum converges. This means that $a > 0$ as long as $N < \infty$.

We proceed to show that $b > 0$ [which is defined in (25)]. We note that $b > q_N$, the probability of N agents absorbing at zero. To see that this is true, consider the last N agents and some order of games for them to play against each other. By the induction hypothesis, this group of N players have positive probability of always playing a doubt. But this group and these games are interspersed with matches between some agent in the group of N and the first agent. We have conditioned on both agents playing doubt in all games until the current game, which includes those involving the first agent. This implies that the beliefs of agents who were chosen for games against the first agent have a lower belief after this game and play doubt at an even higher probability in their next match than in the original set of games in which there was already a probability $q > 0$ of all agents always doubting.

We have, thus, shown that $a > 0$ and $b > 0$ and, therefore, conclude that also $a \cdot b > 0$. \square

We now state a similar lemma for absorption of N agents around one.

Lemma 2 (Absorption at one for N agents). *Let $\alpha \in (0, 1)$, $\epsilon \in (0, \alpha)$, and suppose $\mathbf{x}(t_0) \in [1 - \epsilon, 1]^N$ for some $t_0 \in \mathbb{N}$, then*

$$\begin{aligned} \mathbb{P}(\mathbf{x}(t) \in [1 - \epsilon, 1]^N, \forall t \geq t_0 \mid \mathbf{x}(t_0) \in [1 - \epsilon, 1]^N) \\ = \mathbb{P}_{[1-\epsilon, 1]^N}(\mathbf{x}(t) \in [1 - \epsilon, 1]^N, \forall t \geq t_0) > 0. \end{aligned} \quad (26)$$

The proof is similar to the proof of Lemma 1. The difference is that instead of playing doubt, the agents are required to play trust indefinitely. This happens at probability $F(x)$ rather than $\bar{F}(x)$. Furthermore, the agent belief at the start of the l th round after t_0 is $1 - (1 - x(t_0))\bar{\alpha}^{l-1}$ rather than $x(t_0)\bar{\alpha}^l$.

Intuitively, our first two lemmas imply that a population that is within an ϵ -ball around zero or one can remain there. This means a population believing that everyone is 100% (or almost 100%) trustworthy or untrustworthy can retain that belief forever.

B. Reaching the corners is possible

Our next result proves sufficient conditions for the population to reach the ϵ -ball around zero with positive probability. For legibility, we define

$$I_N := [0, 1]^N \setminus ([0, \epsilon]^N \cup [1 - \epsilon, 1]^N), \quad (27)$$

which we call the *interior*.

Lemma 3 (Population of N agents reaches zero with positive probability). *Let $\alpha \in (0, 1)$ and $\mathbf{x}(t_0) \in I_N$ for some $t_0 \in \mathbb{N}$ and*

$\epsilon \in (0, \alpha)$, then

$$\mathbb{P}_{I_N}(\exists t_1 \in \mathbb{N} : \mathbf{x}(t_1) \in [0, \epsilon]^N, t_1 > t_0) > 0.$$

Proof. We construct a path from I_N , which depends on agents always playing doubt against one another.

We split this into two cases; in the first, we show that there is a path from $\mathbf{x}(t_0) \in (\epsilon, 1 - \epsilon)^N$ to $(0, \epsilon)^N$ of positive probability. In the second, we show that there is a path from $\mathbf{x}(t_0) \in I_N$ with some number $0 < h < N$ of agents with belief $x(t_0) \geq 1 - \epsilon$ to $(\epsilon, 1 - \epsilon)^N$ of positive probability and, therefore, also a path to $(0, \epsilon)^N$.

Case 1: Suppose all agents have belief $x_i(t_0) < 1 - \epsilon$. From now until round m , agent $2k - 1$ is matched to play against agent $2k$ where $k = r \bmod N/2^{60}$ where $r = t - t_0$. The probability of each round of this matching is given by $1/(N(N - 1))$. The probability of this pattern of matching for m rounds is then given by

$$p_m := \prod_{r=1}^m \frac{1}{N} \frac{1}{N - 1} = \frac{1}{N^m(N - 1)^m}.$$

Note that each pair of agents ($2k - 1$ and $2k$) is independent of the other agents for these m rounds, in which they each play $2m/N$ games. Let $t_{i,l}$ for $l = 1, 2, \dots, 2m/N$ index the time of the round in which agent i plays their l th game (after t_0), then supposing all agents play D in all rounds until round m , then agent i 's belief follows:

$$x_i(t_{i,l}) = \bar{\alpha}^l x_i(t_0) \quad \text{for all } l = 1, 2, \dots, 2m/N.$$

Let $\kappa_i \in \mathbb{N}$ be the minimum number of games agent i has to play (both players always playing doubt D) for their belief to be distance ϵ from 0. Specifically, κ_i is the least value that satisfies

$$\bar{\alpha}^{\kappa_i} x_i(t_0) < \epsilon. \quad (28)$$

By dividing through by $x_i(t_0) > 0$ and taking the logarithm, we see that

$$\kappa_i = \left\lceil \frac{\log(\epsilon/x_i(t_0))}{\log \bar{\alpha}} \right\rceil. \quad (29)$$

This is finite for all $x_i(t_0) \leq 1 - \epsilon$, $\epsilon > 0$ and $\alpha < 1$ and is maximized at $x_i(t_0) = 1 - \epsilon$.

For $\mathbf{x}(t) \in (0, \epsilon)^N$, we need $x_i(t) < \epsilon$ to be true for all agents $i = 1, 2, \dots, N$. We define m to be the least value m such that each agent has played enough games to be less than distance ϵ from zero,

$$m := \min\{m \in \mathbb{N} : 2m/N \geq \kappa_i, \forall i = 1, 2, \dots, N\}.$$

By rearrangement, we have

$$m = \max_{i=1,2,\dots,N} \left\lceil \frac{N \cdot \kappa_i}{2} \right\rceil, \quad (30)$$

which is finite because κ_i is finite. This gives us our first intermediate result: The probability of the matching we have created is, thus, strictly greater than zero. We call this $p_m > 0$.

We now turn to the other requirement of this path to $(0, \epsilon)^N$, which is that all agents play D in each of their $2m/N < \infty$ games. As noted before, each pair of agents interacts exclusively with one another and therefore, we focus on one such pair and the rounds in

which they play. For pair $k = 1, 2, \dots, N/2$ denote the probability that there exists a round t_k for which $x_{2k-1}(t_k), x_{2k}(t_k) < \epsilon$:

$$\mathbb{P}_{I_N}(\exists t_k \in \mathbb{N} : x_{2k-1}(t_k), x_{2k}(t_k) < \epsilon) \geq \prod_{l=0}^{2m/N} \bar{F}(\bar{\alpha}^l x_{2k-1}(t_0)) \bar{F}(\bar{\alpha}^l x_{2k}(t_0)). \tag{31}$$

Let $\tilde{x}_{k,0} := \max\{x_{2k-1}, x_{2k}\}$ and subsequently $\tilde{x}_{k,l} = \bar{\alpha}^l \tilde{x}_{k,0}$. We use this to bound the beliefs of the k -th pair in their l -th interaction. Then we can bound the right hand side:

$$\mathbb{P}_{I_N}(\exists t_k \in \mathbb{N} : x_{2k-1}(t_k), x_{2k}(t_k) < \epsilon) \geq \prod_{l=0}^{2m/N} (\bar{F}(\tilde{x}_{k,l}))^2, \tag{32}$$

because $\tilde{x}_{k,l} \geq x_{2k-1}(t_0)\bar{\alpha}^l, x_{2k}(t_0)\bar{\alpha}^l$ for all $l \in \mathbb{N}$ and because $F(x) \geq F(y)$ as long as $x \geq y$ by virtue of being a distribution function. The terms in the product are all strictly greater than 0 and because it is a finite product, we know that its result is also strictly greater than 0 for all $k = 1, 2, \dots, N/2$. Finally, we have $N/2$ of these pairs and so the probability of reaching the corner is the probability of the matching times the probability of the individual games proceeding as described,

$$\mathbb{P}_{I_N}(\exists t_1 \in \mathbb{N} : \mathbf{x}(t_1) \in [0, \epsilon]^N, t_1 > t_0) > p_m \cdot \mathbb{P}_{I_N}(\exists t_k \in \mathbb{N} : x_{2k-1}(t_k), x_{2k}(t_k) < \epsilon)^{N/2} > 0. \tag{33}$$

This proves the statement in case 1 where all the agent beliefs were $x_i(t_0) \leq 1 - \epsilon$ for all $i = 1, 2, \dots, N$.

Case 2: Alternatively, there are agents who have belief $x_i(t_0) > 1 - \epsilon$. As a worst case scenario, suppose $x_1(t_0) = 1 - \epsilon$ and $x_j(t_0) = 1$ for $j = 2, \dots, N$. Here, we have chosen the first agent to be the agent with belief $x(t_0) \leq 1 - \epsilon$ without loss of generality as their naming convention plays no role. We now construct a finite path of positive probability to get from this state to the state we assumed at the start of the proof [$x_i(t_0) < 1 - \epsilon$ for all $i = 1, 2, \dots, N$].

For $k = 1, 2, \dots, N - 1$, we specify the games between I and J as follows:

$$(I(t), J(t))_{t=t_0+1, \dots, t_0+2(N-1)}^* = \begin{cases} I(t_0 + 2k - 1) = I(t_0 + 2k) \in \{1, \dots, k\} \\ J(t_0 + 2k - 1) = J(t_0 + 2k) = k + 1. \end{cases}$$

The probability of such a sequence of matches is

$$p'_m := \mathbb{P}((I(t), J(t))^*) = \prod_{k=1}^{N-1} \binom{k}{N} \binom{1}{N} \left(\frac{1}{N-1}\right)^2 > 0. \tag{34}$$

The above is a finite product of positive numbers and so is greater than zero. The general form of this sequence is that some agent from the set $\{1, \dots, k\}$ plays against the first agent that is not in the set twice. In the first of these two games, player I plays doubt in both rounds, while player J plays trust in the first round and doubt in the second round. In this way, after the two games, both agents have $x < 1 - \epsilon$.

To see that this is the case consider agent I whose belief starts at $x(t_0) \leq 1 - \epsilon$. After playing against T , their belief updates to $x(t_0 + 1) \leq 1 - \epsilon\bar{\alpha} < 1$. Subsequently, they play against D , and so in the following round, they hold belief $x(t_0 + 2) \leq 1 - \alpha - \epsilon\bar{\alpha}^2 < 1 - \alpha < 1 - \epsilon$. Now, we consider the agent J whose belief starts at $x(t_0) \leq 1$. After playing against D twice their belief is bounded by $x(t_0 + 2) \leq \bar{\alpha}^2 < 1 - \alpha < 1 - \epsilon$.

For one such pair of games, define B the event of the sequence of actions (starting in round $2k + 1$),

$$B := \{A_1(t + 2k - 1) = D\} \cap \{A_{-1}(t + 2k - 1) = T\} \cap \{A(t + 2k) = D\}.$$

We call p_{2g} , the probability of both players acting according to the event B , and we note that it is bounded,

$$p_{2g} := \mathbb{P}_{I_N}(B) \geq \bar{F}(1 - \epsilon) \cdot 1 \cdot \bar{F}(1 - \epsilon\bar{\alpha}) \cdot \bar{F}\bar{\alpha}. \tag{35}$$

The right hand side of (35) is strictly greater than zero because by Assumption 3, F is only zero at zero, and therefore, we have the product of four numbers, all greater than zero. Thus, $p_{2g} > 0$.

The probability then of getting into case 1 is given by the probability of the matching $(I, J)^*$ multiplied by the probability of all $(N - 1)$ pairs of games going as planned, which are independent and the probability is

$$p'_m \cdot p_{2g}^{N-1} > 0. \tag{36}$$

At time $t_0 + 2(N - 1) + 1$, all the agents' belief is $x < 1 - \epsilon$. Proceeding as in case 1, we know that there is a path of positive probability to $(0, \epsilon)^N$. \square

We now state a similar lemma for the probability of the population reaching the ϵ -ball around 1 from the interior.

Lemma 4 (Population of N agents reaches one with positive probability). Let $\alpha \in (0, 1)$ and $\mathbf{x}(t_0) \in I_N$ for some $t_0 \in \mathbb{N}$ and $\epsilon \in (0, \alpha)$, then

$$\mathbb{P}_{I_N}(\exists t_1 \in \mathbb{N} : \mathbf{x}(t_1) \in [1 - \epsilon, 1]^N, t_1 > t_0) > 0.$$

The proof is similar to that of Lemma 3. The differences are akin to the differences between Lemma 1 and Lemma 2. Additionally, instead of looking for a least κ_i that satisfies (28), we look for a least κ_i that satisfies $1 - (1 - x(t_0))\bar{\alpha}^{\kappa_i} > 1 - \epsilon$. This can, however, be translated into (28). We now state the main analytical result of the paper. By the dynamics we describe, a population of any finite size N is guaranteed to converge, in belief as well as behavior, to one of the corners $\mathbf{0}$ or $\mathbf{1}$ of the state space $[0, 1]^N$.

Our second pair of lemmas tell us that a population can end up believing that the rest of the population is 100% trustworthy (or untrustworthy). Thus a natural reinforcement of beliefs regarding trustworthiness can take place and result in complete trust among the population or a complete lack thereof.

C. Proof of the main theorem

We now present the proof of the main theorem of our paper, which states that a population of agents will converge with probability one at either the always-trust or the always-doubt corner.

Proof. Our objective is to show that if $\alpha > \epsilon$, and $F^{(n_1)}(0)$, $F^{(n_2)}(1) \in [0, \infty)$ for some $n_1, n_2 < \infty$, then

$$\mathbb{P}(\exists t_0 : \mathbf{x}(t) \in [0, \epsilon]^N \cup [1 - \epsilon, 1]^N, \forall t > t_0) = 1.$$

By Lemmas 3 and 4, we know that the process may reach $A := [0, \epsilon]^N \cup [1 - \epsilon, 1]^N$ from $[0, 1] \setminus A$ in s (a finite number of) rounds with positive probability.

By Lemmas 1 and 2, we know that the belief vector, upon reaching A , has a positive probability of being absorbed there. As such, we define the probability of reaching and being absorbed in A in $s + 1$ rounds as

$$p := \mathbb{P}(\mathbf{x}(t) \in A, \forall t \geq t_0 + s + 1),$$

and note that $p > 0$. Define τ_1 as the time of the last entry into A ,

$$\tau_1 := \min\{\tau : \mathbf{x}(t) \in A, \forall t \geq \tau\} \in \mathbb{N} \cup \{\infty\}.$$

Thus, after $s + 1$ rounds (supposing we can read the future about possible absorption), one of the three things will have occurred:

1. The belief vector $\mathbf{x}(t_0 + s + 1)$ is absorbed in A , and so, $t_0 + s + 1 \geq \tau_1$.
2. The belief vector $\mathbf{x}(t_0 + s + 1)$ is not absorbed in A (thus, $t_0 + s + 1 < \tau_1$) and is somewhere in $[0, 1]^N \setminus A$.
3. The belief vector $\mathbf{x}(t_0 + s + 1)$ is in A (thus, $t_0 + s + 1 < \tau_1$) but has not been absorbed.

If (3) happens, then at some point, the belief will have to again be in $[0, 1]^N \setminus A$ (with probability 1 as this is implied by not being absorbed). Once the process is again in $[0, 1]^N \setminus A$ either directly in case of (2) and after some finite time in case (3), we reset the clock and note that after another $s + 1$ rounds, the process is either absorbed in A (with probability p) or not (with probability $1 - p$). We, therefore, know that the probability of the belief vector \mathbf{x} not being absorbed in A after $n \in \mathbb{N}$ such meta-experiments is given by $(1 - p)^n$ and the probability of never being absorbed is

$$\mathbb{P}(\nexists t_0 : \mathbf{x}(t) \in A, \forall t \geq t_0) = \lim_{n \rightarrow \infty} (1 - p)^n = 0.$$

Thus, the probability of the complement is $\mathbb{P}(\exists t_0 : \mathbf{x}(t) \in A, \forall t \geq t_0) = 1$.

In this section, we showed that it is certain that the process ends in one of the corners in the long run. Note that it might be possible to prove this result by carefully checking whether the dynamics of this model satisfy the conditions in Norman.⁶⁸ We do not believe this to be a straightforward task. The natural next question is: With what probability does the system absorb in the always-trust corner? Define the probability of (eventual) absorption in the always-trust corner $\mathbf{x} \in [1 - \epsilon, 1]^N$ as

$$p_T(\mathbf{x}(t_0)) := \mathbb{P}(\exists t_1 : \mathbf{x}(t) \in [1 - \epsilon, 1]^N, \forall t \geq t_1).$$

Similarly, we define the probability of absorption in the always-doubt corner,

$$p_D(\mathbf{x}(t_0)) := \mathbb{P}(\exists t_1 : \mathbf{x}(t) \in [0, \epsilon]^N, \forall t \geq t_1).$$

Because absorption in one of the two corners is guaranteed by Theorem 1, we have the following relationship for all $\mathbf{x} \in [0, 1]^N$:

$$p_D(\mathbf{x}) = 1 - p_T(\mathbf{x}).$$




FIG. 2. The expected value of $h(\mathbf{x}) - \mathbf{x}$ for the function $F(x) = x + 2x(x - 1/2)(x - 1)$. This figure highlights the at first counterintuitive nature of the result in Theorem 1. The expectation of the dynamics would indicate that the process absorbs at $\mathbf{x} = (0.5, 0.5)$, but we know by Theorem 1 that the process absorbs with probability one at one of the two corners $\mathbf{x} = (0, 0)$ or $(1, 1)$.

By this result, we are able to focus our investigation into the dichotomy of absorption at $\mathbf{1}$ or at $\mathbf{0}$ and not miss dynamics in which the process reaches a steady state in the interior of the state space.

Remark 2. The main result of this section, Theorem 1 also holds true for any finite connected population structure determining agent pairings. Absorption of the process \mathbf{x} at $\mathbf{0}$ or $\mathbf{1}$ is guaranteed in the long-term.

The proofs for the corresponding results of Lemmas 1–4 would need to be adjusted by accounting for the restrictions on which agents may interact with one another. With some care on the order of agent pairings, as long as the network of agents is connected, it should be possible to construct a sequence of pairings, which leads to reaching the ϵ -balls at zero and one. The same holds true for absorption.

IV. ILLUSTRATION FOR TWO AGENTS

In this section, we illustrate the model by means of exploring the two agent case. Suppose that $N = 2$, then we have two agents who are matched to play against one another in all rounds $t = 1, 2, \dots$. Note that the agents assignment to I or J are merely a matter of notation. Thus, we can assume that agent 1 is always assigned $k = 1$ and agent 2 is assigned $k = -1$ and refer to them by their k assignment for the remainder of this section.

We have the following stochastic dynamical system:

$$\begin{bmatrix} x_1(t+1) \\ x_{-1}(t+1) \end{bmatrix} = h(\mathbf{x}(t)) + \xi_t, \tag{37}$$

where ξ_t is defined as the vector of noise, which accounts for the randomness,

$$\xi_t := \begin{bmatrix} \bar{\alpha}x_1(t) + \alpha \mathbb{1}_{\{A_{-1}(t)=T\}} \\ \bar{\alpha}x_{-1}(t) + \alpha \mathbb{1}_{\{A_1(t)=T\}} \end{bmatrix}, \quad (38)$$

and we define $h(\cdot)$ to be the expected evolution of the system. We may calculate $h(x)$ explicitly in terms of the cdf F ,

$$h(x(t)) = \begin{bmatrix} \bar{\alpha}x_1(t) + \alpha F(x_{-1}(t)) \\ \bar{\alpha}x_{-1}(t) + \alpha F(x_1(t)) \end{bmatrix}.$$

We note that any points with $x_1 = F(x_{-1})$ and $x_{-1} = F(x_1)$ map to themselves in expectation. Furthermore, because $F(x) \geq F(y)$ whenever $x \geq y$ (because F is a distribution function), the points that map to themselves in expectation must satisfy

$$x_1 = x_{-1}, \quad (39)$$

$$x_1 = F(x_1). \quad (40)$$

For the function $F(x) = x + 2x(x - 1/2)(x - 1)$, we plot the vector field of the expected change: $h(x) - x$ in Fig. 2. Looking at the regions $(0, 0.4)^2$ and $(0.6, 1)^2$ in Fig. 2, we might guess that when $F(x) = x + 2x(x - 1/2)(x - 1)$ convergence will be to the center of the domain: $x = (0.5, 0.5)$. However, by Theorem 1, we know that the process must converge in one of the two corners in the long run.

Heuristically, this can be explained by the progression of agent beliefs around zero and one vs around 0.5. There are paths by which the agents' beliefs can keep getting closer to zero and one, respectively. Meanwhile, at 0.5, should the agents' beliefs get arbitrarily close, their next belief will be roughly $|0.5\alpha|$ away from 0.5.


In short, knowing the result of Theorem 1 and comparing this with Fig. 2, we realize that we should be careful not to assume that the asymptotic dynamics will be dictated by their expectation.

V. SIMULATION


We analyze the effect of F and α on p_T and the time until first entry to A by means of simulation. Because it requires substantially less simulated time steps to simulate the process with $N = 2$ agents, we begin in this setting and later include a numerical simulation to verify that the results are qualitatively similar to the case when $N > 2$.

We parameterize F by $r \in [0.5, 2]$ and take only functions of the form $F(x) = x^r$. In particular, note that $F(x) = \sqrt{x}$ and $F(x) = x^2$ are included in this parameterization. The two agents are initialized with beliefs chosen uniformly at random over the belief space $(0, 1)$. Thus, $x(0) \sim \mathcal{U}_{[0,1]}^2$. In order to determine how many time steps ought to be simulated to allow process to be absorbed in one of the two steady states, we run a preliminary simulation for the values $\alpha = \{0.01, 0.05, 0.25, 0.5\}$ and $r = \{0.5, 1, 2\}$. We run 1000 iterations for 10 000 times steps keeping track of the average Manhattan distance to the nearest steady state over time. The resulting 95% confidence bounds of the average Manhattan distance to A is depicted in Figs. 3(a)–3(d).


In Fig. 3(a), we see that assuming convergence at some time $t < 10\,000$ when $r = 1$ and $\alpha \leq 0.01$ is devious and should be avoided. For $\alpha = 0.05, 0.25$, and 0.5 , we see that 5000, 250, and 100 rounds, respectively, are enough to be 95% confident that the




(a) $\alpha = 0.01$



(b) $\alpha = 0.05$



(c) $\alpha = 0.25$



(d) $\alpha = 0.50$

FIG. 3. Average minimum Manhattan distance to A over time for various α and r .




FIG. 5. Heatmaps of $p_T(\mathbf{x}(0))$ and the time it takes to be absorbed at either corner. Increasing the number of agents in the population dampens the effect of the model learning rate α on the outcome or sharpening the effect of the payoff distribution parameter r .

process has indeed been absorbed in a steady state. For the simulation study, we run the simulation for 1000 time steps when $\alpha \geq 0.25$ and for 10 000 time steps for $\alpha < 0.25$. In particular, we run simulations for $\alpha \in \{0.05, 0.1, 0.2, 0.3, 0.4, 0.5\}$ for values of $r \in \{0.5, 0.6, 0.7, 0.8, 0.9, 0.95, 1, 1.05, 1.1, 1.25, 1.40, 1.55, 2\}$.

We show the probability of being absorbed in the always-trust corner of the state space in Fig. 4(a). We notice that increasing α has the effect of smoothing the transition between being absorbed in (0, 0) and in (1, 1). Conversely, decreasing α results in a sharper transition between being absorbed in (0, 0) with probability one when $r > 1$ and absorbed in (1, 1) with probability one when $r < 1$. The following conjecture elaborates on this.

Conjecture 1 (Phase transition at $r = 1$ as $\alpha \rightarrow 0$). As $\alpha \rightarrow 0$, $p_T(\mathbf{x}) = 1$ when $r < 1$ and $p_T(\mathbf{x}) = 0$ when $r > 1$ for all $\mathbf{x} \in [0, 1]^N \setminus \{\mathbf{0}^N, \mathbf{1}^N\}$.

The machinery used in the proofs requires $\epsilon < \alpha$. This provides clear criteria for events that lead the population to “jump out” of the ϵ -balls around zero and one. It also ensures that the population can converge to an ϵ -ball around zero or one in finite time [see division by $\log(\bar{\alpha}) = \log(1 - \alpha)$ in (29), for example]. New techniques would, thus, have to be applied in order to prove Conjecture 1. Intuitively however, as the learning rate α goes to zero, the size of the steps agents make when they adjust their belief shrink. This seems to imply that an infinite number of steps might be needed, but it is only their direction that matters and so movement toward the expectation (to zero when $r > 1$ and toward one when $r < 1$) will eventually come to pass.

Phase transitions are quite common in evolutionary game theory, with some recent results in Refs. 35, 38, and 41. Duong and Han³⁵ also conjecture a phase transition (regarding the expected

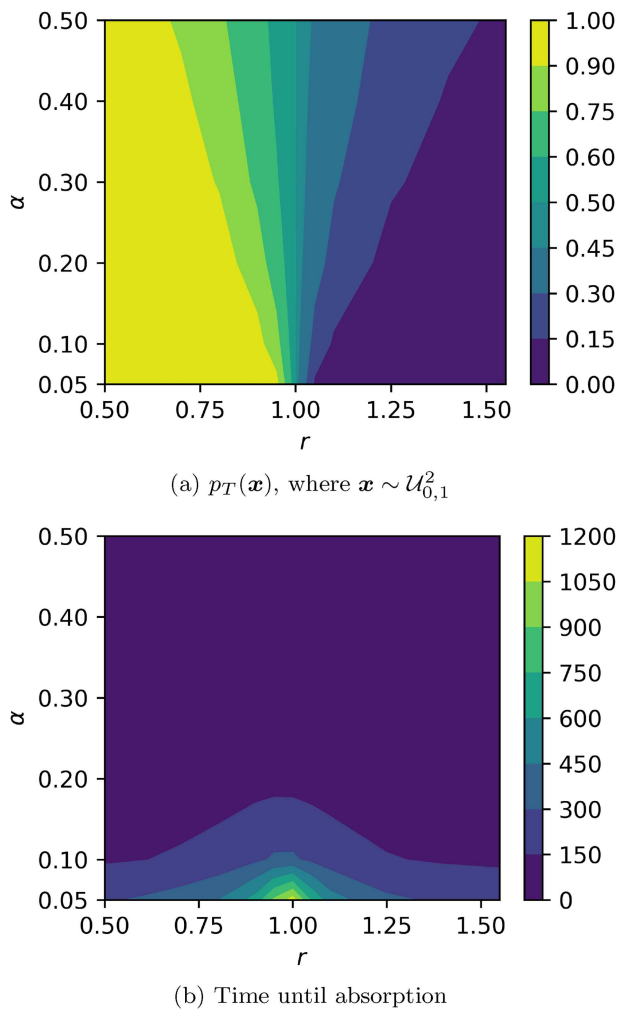


FIG. 4. Heatmaps of $p_T(\mathbf{x}(0))$ and the time it takes to be absorbed at either corner. We see the transition between being absorbed in $(1, 1)$ with probability 1 when $r < 1$ and probability zero when $r > 1$. A greater α has the effect of smoothing this transition. We also see that increasing α has a hastening effect on the time to absorption.

number of equilibria) whose proof is beyond the standard techniques for phase transitions in the replicator dynamics. Zheng *et al.*³⁸ and Zeng *et al.*⁴¹ present phase transitions for imitation learning dynamics, which they observe in simulations. Though beyond the scope of this paper, Conjecture 1 could be investigated in future work by means of an extensive simulation study.

In Fig. 4(b), we show the time average time until the simulation entered A for the first time. By choosing $\epsilon = 10^{-5}$, we suggest that the time it takes to enter $A = [0, \epsilon]^2 \cup [1 - \epsilon, 1]^2$ is proportional to the time it takes to also be absorbed in A . Supposing this is the case, then we see that a greater α has the effect of speeding up the dynamics. This is sensible as the step sizes made by the process are bigger for a bigger α . We also see that the closer r is to 1, the longer the

dynamics take. This is explainable by the fact that with a greater r , while the process is still far from A , the probability of steps toward the center is still fairly likely. In short, smaller $|1 - r|$ implies that the process spends more time in the center of the state space, while lower α means that the process is moving to the corners at a slower pace.

In order to show how the behavior changes as the population size is increased, we present simulations for $N = 5$ and $N = 10$. The resulting probability of absorption at $\mathbf{1}^N$ is depicted in a heatmap-format in Figs. 5(a)–5(b). For $N = 5$, in Fig. 5(a), we already see that the effect of α observed for $N = 2$ is less pronounced, though still present. Furthermore, for $N = 10$ in Fig. 5(b), we see that α almost plays no role in terms of the relative absorption probabilities. Thus, as N increases, we suspect that the steady state in which the process will be absorbed depends more on the cdf F and less on the learning dynamics of the individual agents.

The time to absorption for the simulation runs with $N = 5$ and $N = 10$ is shown in Figs. 5(c) and 5(d), respectively. We see that, qualitatively, the interplay between α and r on the time to absorption is the same. That is, bigger α and $|1 - r|$ speed up the dynamics.

VI. CONCLUSION

We modeled the problem of trusting strangers in society as game in which agents of a population are randomly matched, two per round, and tasked with a coordination game with random payoffs. These agents are endowed with a learning procedure by which they update their belief on the probability that a random stranger would trust using the exponential moving average of their passed observations. We have shown for that for any finite population of size N , with mild conditions on the cdf of the payoff parameters, and a constant learning rate (or memory) $\alpha \in (0, 1)$, the process is absorbed with probability 1 at one of the two steady states: always-trust and always-doubt. This result is not immediately obvious because looking at the expected change of the process, there are F for which it would seem that all agents believing $\mathbf{x} = [0.5]^N$ is an attracting steady state.

Furthermore, we analyzed the effect of F , α , and N by simulating the model with parameterized $F(x) = x^r$ for $r \in [0.5, 2]$ and $\alpha \in [0.05, 0.5]$. The results of this simulation show that

- decreasing α exaggerates the effect r and
- increasing N decreases the effect of α .

We conjecture that as $\alpha \rightarrow 0$, there will emerge a sharp phase transition at $r = 1$ such that the process is absorbed in the always-trust steady state with probability 1 when $r < 1$ and with probability 0 when $r > 1$.

A similar phase transition might occur as the population size is increased. A broad takeaway of this model is that differences in trust between populations might simply boil down to chance; the process by which a population learns on which action to coordinate is random and is not necessarily fully determined by the nature of the game (F), the size of the population (N), or the rate of learning (α). However, in the case of large populations, our simulations suggest that the context of the interaction (distributions of the payoff parameters) will play a role in determining whether or not they end up with long-term trust or long-term doubting. In terms of the

numerical simulation, the probability of converging to the always-trust steady state decreases (increases) as the population grows for $r > 1$ ($r < 1$). This implies that the effect of increasing population is not determined by the type of game (like for the one-shot N -player public goods and prisoner's dilemma games shown in Ref. 61) but really the specific payoff structure. For instance, our model corresponds to an all-or-nothing version of the 2 player public goods game when F is a step function with one step from zero to one. Players either pay one (when $A_k = T$) or zero (when $A_k = D$), while the reward $\gamma \in \mathbb{R}$ is only paid out if both players payed the cost. Then, whether or not a larger population benefits one or the other steady state depends not on the fact that it is a public goods game but specifically where the step function steps from zero to one. Specifically, the location of this step corresponds to $1/\gamma$.

As a possible answer to our research question, we see that societies may be low or high trust largely due to the context of the interactions within them. The context being the distributions of the payoffs in the “coordination game” the people are playing with one another. Adjusting the payoffs of the social interactions might, thus, be the most effective way to foster more trust in a population. Chance plays a bigger role in smaller groups than with large ones. Future research could be to investigate whether a large population could be modeled by a loose grouping of smaller cliques and that the average trusting behavior of the population is an aggregation of these cliques. In this case, chance may indeed play a big role in determining the amount of trust on a clique level, which indirectly effects the trust in the population as a whole.

The model we consider assumes a population structure in which all agents are connected with one another. Although the result of Theorem 1 should also hold for any other network topology as long as this is connected, it might be interesting to look for the possibility of pseudo-steady states in networks with strong community structures. It could also be interesting to characterize mapping between the structure of the graph defining the interactions and the rate of convergence to a steady state. Other future work could focus on extending the convergence result (or proving its opposite) in the case of Bayesian learning or another simple learning rule that the agents might employ.

Our model is such that if there exists even one agent who always trusts regardless of the actions taken by the rest of the population, the population will converge to the trust action with probability one (because convergence to always doubt is impossible). This is in contrast to Wang and Sun¹⁶ in whose model the presence of “zealous” always cooperators does not always promote cooperation. However, in our model, the converse is also true: One constant doubter forces the always-doubt equilibrium. Studying the transient dynamics of our model with both constant trusters and constant doubters might lead to results with oscillations between trusting and doubting as in Ref. 62.

Another line of future work could do away with the assumption of selfish rationality and instead have agents using some form of moral preferences showcased in Refs. 31 and 63. For example, agents who are interested in the welfare of both players in the game. A facet of trust we have not explored here has to do with believing what has been said by others. Recently, the evolution of honesty in the sender–receiver game has been studied under imitation of better-performing strategies by means of the Monte Carlo method.^{64–66} In

particular, agents in these studies copy the behavior of other agents at probability proportional to the difference between their payoffs. It may be fruitful to adapt the approach we have taken here to study the evolution of honesty in the sender–receiver game under experience based learning instead of imitation. In this case, it may be necessary for agents to have two beliefs: x (likelihood of being believed) and y (likelihood of others being honest). In particular, it may be interested to investigate under which conditions the golden rule (lie and disbelieve or be honest and believe) emerges.

ACKNOWLEDGMENTS

This research was supported by the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie Grant Agreement No. 945045 and by the NWO Gravitation Project NETWORKS under Grant No. 024.002.003.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Benedikt V. Meylahn: Conceptualization (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Software (equal); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal). **Arnoud V. den Boer:** Conceptualization (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Supervision (equal); Writing – original draft (equal); Writing – review & editing (equal). **Michel Mandjes:** Conceptualization (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Supervision (equal); Writing – original draft (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

- ¹E. Ortiz-Ospina and M. Roser, *Trust* (Our World in Data, 2016).
- ²T. U. Thomsen and E. B. Sørensen, “The road to motherhood,” *Advert. Soc. Rev.* 7(4), 1–10 (2006).
- ³J. Berg, J. Dickhaut, and K. McCabe, “Trust, reciprocity, and social history,” *Games Econ. Behav.* 10, 122–142 (1995).
- ⁴A. Kumar, V. Capraro, and M. Perc, “The evolution of trust and trustworthiness,” *J. R. Soc. Interface* 17(169), 20200491 (2020).
- ⁵K. Sun, Y. Liu, X. Chen, and A. Szolnoki, “Evolution of trust in a hierarchical population with punishing investors,” *Chaos, Solitons Fractals* 162, 112413 (2022).
- ⁶M. Feng, X. Li, D. Zhao, and C. Xia, “Evolutionary dynamics with the second-order reputation in the networked N -player trust game,” *Chaos, Solitons Fractals* 175, 114042 (2023).
- ⁷Y. Liu, L. Liu, R. Guo, and L. Zhang, “ N -player repeated evolutionary trust game under government management,” *Chaos, Solitons Fractals* 175, 114031 (2023).
- ⁸C. Wang, “Evolution of trust in structured populations,” *Appl. Math. Comput.* 471, 128595 (2024).
- ⁹Fictitious play describes the agent learning paradigm in which each agent assumes that they are playing against a fixed mixed strategy profile. The agents

use historical play to try and learn the profile they believe to be playing against and act myopically optimally against it.⁶⁸ Our agents' actions are akin to fictitious play as they base these myopically on their belief. However, their belief is updated as if their opponents strategy can change.

¹⁰A. P. Burger, M. Van der Merwe, and J. H. van Vuuren, "The evolutionary spatial prisoner's dilemma on a cycle," *ORiON* **29**(1), 1–16 (2013).

¹¹I. S. Lim and P. Wittek, "Satisfied-defect, unsatisfied-cooperate: An evolutionary dynamics of cooperation led by aspiration," *Phys. Rev. E* **98**, 062113 (2018).

¹²I. S. Lim, "Stochastic evolutionary dynamics of trust games with asymmetric parameters," *Phys. Rev. E* **102**, 062419 (2020).

¹³B. V. Meylahn and J. H. van Vuuren, "The evolutionary spatial snowdrift game on a cycle: An asymptotic analysis," *Open J. Discrete Appl. Math.* **4**(3), 36–59 (2021).

¹⁴P. Godara, T. D. Aléman, and S. Herminghaus, "Bounded rational agents playing a public goods game," *Phys. Rev. E* **105**, 024114 (2022).

¹⁵Z. Yang and L. Zhang, "Random migration with tie retention promotes cooperation in the prisoner's dilemma game," *Chaos* **33**(4), 043126 (2023).

¹⁶C. Wang and C. Sun, "Zealous cooperation does not always promote cooperation in public goods games," *Chaos* **33**(6), 063111 (2023).

¹⁷D. Fudenberg and D. M. Kreps, "Learning mixed equilibria," *Games Econ. Behav.* **5**, 320–367 (1993).

¹⁸J. Hofbauer and W. H. Sandholm, "On global convergence of stochastic fictitious play," *Econometrica* **70**(6), 2265–2294 (2002).

¹⁹Y. M. Kaniovski and H. P. Young, "Learning dynamics in games with stochastic perturbations," *Games Econ. Behav.* **11**, 330–363 (1995).

²⁰K. Tuyls and A. Nowé, "Evolutionary game theory and multi-agent reinforcement learning," *Knowl. Eng. Rev.* **20**(1), 63–90 (2006).

²¹R. A. Laird, "Static cooperator-defector patterns in models of the snowdrift game played on cycle graphs," *Phys. Rev. E* **88**, 012105 (2013).

²²L. Baudin and R. Laraki, "Smooth fictitious play in stochastic games with perturbed payoffs and unknown transitions," in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., 2022), Vol. 35, pp. 20243–20256.

²³P. Godara and S. Herminghaus, "Bounded learning and planning in public goods games," *Phys. Rev. E* **107**, 054140 (2023).

²⁴B. V. Meylahn, A. V. den Boer, and M. R. H. Mandjes, "Trusting: Alone and together," *J. Math. Sociol.* [arXiv:2303.01921](https://arxiv.org/abs/2303.01921) (2023).

²⁵Y. Lu, Y. Wang, Y. Liu, J. Chen, L. Shi, and J. Park, "Reinforcement learning relieves the vaccination dilemma," *Chaos* **33**(7), 073110 (2023).

²⁶P. Zhu, Z. Cao, C. Liu, C. Chu, and Z. Wang, "Co-evolution of synchronization and cooperation with multi-agent Q-learning," *Chaos* **33**(3), 033128 (2023).

²⁷H. P. Young, "The evolution of conventions," *Econometrica* **60**(1), 57–84 (1993).

²⁸H. Xie and Y.-J. Lee, "Social norms and trust among strangers," *Games Econ. Behav.* **76**, 548–555 (2012).

²⁹E. Bilancini and L. Boncinelli, "The evolution of conventions under condition-dependent mistakes," *Econ. Theory* **69**, 497–521 (2020).

³⁰M. O. Sayin, F. Parise, and A. Ozdaglar, "Fictitious play in zero-sum stochastic games," *SIAM J. Control Optim.* **60**(4), 2095–2114 (2022).

³¹J. A. Joireman, G. P. Shelley, P. D. Teta, J. Wilding, and D. Michael Kuhlman, "Computer simulation of social value orientation: Vitality, satisfaction, and emergent game structures," in *Frontiers in Social Dilemmas Research*, edited by W. B. G. Liebrand and D. M. Messick (Springer, Berlin, 1996), pp. 289–310.

³²E. Ert, I. Erev, and A. E. Roth, "A choice prediction competition for social preferences in simple extensive form games: An introduction," *Games* **2**, 257–276 (2011).

³³J. Berg and A. Engel, "Matrix games, mixed strategies, and statistical mechanics," *Phys. Rev. Lett.* **81**, 4999–5002 (1998).

³⁴T. A. Han, A. Traulsen, and C. S. Gokhale, "On equilibrium properties of evolutionary multi-player games with random payoff matrices," *Theor. Popul. Biol.* **81**(4), 264–272 (2012).

³⁵M. H. Duong and T. A. Han, "On equilibrium properties of the replicator-mutator equation in deterministic and random games," *Dyn. Games Appl.* **10**(3), 641–663 (2020).

³⁶S. Karlin and U. Liberman, "Random temporal variation in selection intensities: Case of large population size," *Theor. Popul. Biol.* **6**, 355–382 (1974).

³⁷S. Karlin and U. Liberman, "Random temporal variation in selection intensities: One-locus two-allele model," *J. Math. Biol.* **2**, 1–17 (1975).

³⁸X.-D. Zheng, C. Li, S. Lessard, and Y. Tao, "Evolutionary stability concepts in a stochastic environment," *Phys. Rev. E* **96**, 032414 (2017).

³⁹T.-J. Feng, J. Mei, C. Li, X.-D. Zheng, S. Lessard, and Y. Tao, "Stochastic evolutionary stability in matrix games with random payoffs," *Phys. Rev. E* **105**, 034303 (2022).

⁴⁰M. Perc, "Transition from Gaussian to Levy distributions of stochastic payoff variations in the spatial prisoner's dilemma game," *Phys. Rev. E* **75**, 022101 (2007).

⁴¹Z. Zeng, Q. Li, and M. Feng, "Spatial evolution of cooperation with variable payoffs," *Chaos* **32**(7), 073118 (2022).

⁴²J.-L. Qian, Y.-X. Zhou, and Q.-Y. Hao, "The emergence of cooperative behavior based on random payoff and heterogeneity of concerning social image," *Chaos, Solitons Fractals* **179**, 114405 (2024).

⁴³Y. Sato, E. Akiyama, and J. D. Farmer, "Chaos in learning a simple two-person game," *Proc. Natl. Acad. Sci. U. S. A.* **99**(7), 4748–4751 (2002).

⁴⁴Y. Sato and J. P. Crutchfield, "Coupled replicator equations for the dynamics of learning in multiagent systems," *Phys. Rev. E* **67**, 015206 (2003).

⁴⁵Y. Sato, E. Akiyama, and J. P. Crutchfield, "Stability and diversity in collective adaptation," *Phys. D* **210**(1–2), 21–57 (2005).

⁴⁶S. Banisch and E. Olbrich, "Opinion polarization by learning from social feedback," *J. Math. Sociol.* **43**(2), 76–103 (2019).

⁴⁷J. Nash, "Non-cooperative games," *Ann. Math.* **54**(2), 286–295 (1951).

⁴⁸M. Benaïm and M. W. Hirsch, "Mixed equilibria and dynamical systems arising from fictitious play in perturbed games," *Games Econ. Behav.* **29**, 36–72 (1999).

⁴⁹V. Bala and S. Goyal, "Learning from neighbours," *Rev. Econ. Stud.* **65**(3), 595–621 (1998).

⁵⁰W.-X. Wang, J. Ren, G. Chen, and B.-H. Wang, "Memory-based snowdrift game on networks," *Phys. Rev. E* **74**, 056113 (2006).

⁵¹M. A. Raghunandan and C. A. Subramanian, "Sustaining cooperation on networks: An analytical study based on evolutionary game theory," in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems—Volume 2, AAMAS '12 (Richland, SC)* (International Foundation for Autonomous Agents and Multiagent Systems, 2012), pp. 913–920.

⁵²Y. Li, Y. Min, X. Zhu, and J. Cao, "Partner switching promotes cooperation among myopic agents on a geographical plane," *Phys. Rev. E* **87**, 022823 (2013).

⁵³M. Harel, E. Mossel, P. Strack, and O. Tamuz, "Rational groupthink," *Q. J. Econ.* **136**(1), 621–668 (2021).

⁵⁴H.-W. Lee, C. Cleveland, and A. Szolnoki, "Restoring spatial cooperation with myopic agents in a three-strategy social dilemma," *Appl. Math. Comput.* **458**, 128226 (2023).

⁵⁵S. Zhang and A. J. Yu, "Cheap but clever: Human active learning in a bandit setting," in *Proceedings of the Annual Meeting of the Cognitive Science Society* (Curran Associates, Inc., 2013), Vol. 35.

⁵⁶S. Zhang and A. J. Yu, "Forgetful Bayes and myopic planning: Human learning and decision-making in a bandit setting," in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., 2013), Vol. 26.

⁵⁷H. Robbins and S. Monro, "A stochastic approximation method," *Ann. Math. Stat.* **22**(3), 400–407 (1951).

⁵⁸R. G. Brown, *Smoothing, Forecasting and Prediction of Discrete Time Series* (Prentice-Hall, Englewood Cliffs, NJ, 1963).

⁵⁹E. T. Whittaker and G. N. Watson, *A Course of Modern Analysis*, 3rd ed. (Cambridge University Press, 1920).

⁶⁰This proof handles the case where N is even. To extend the result to all odd $N \geq 3$, a simple extension is possible: The first agent plays a game with the last agent at the start of each iteration step and then joins the pool of $N - 1$ agents, which proceeds as in the case we handle.

⁶¹H. Barcelo and V. Capraro, "Group size effect on cooperation in one-shot social dilemmas," *Sci. Rep.* **5**(1), 7937 (2015).

⁶²A. Szolnoki and M. Perc, "Zealots tame oscillations in the spatial rock-paper-scissors game," *Phys. Rev. E* **93**, 062307 (2016).

⁶³V. Capraro and M. Perc, “Mathematical foundations of moral preferences,” *J. R. Soc. Interface* **18**, 20200880 (2021).

⁶⁴V. Capraro, M. Perc, and D. Vilone, “The evolution of lying in well-mixed populations,” *J. R. Soc. Interface* **16**, 20190211 (2019).

⁶⁵V. Capraro, M. Perc, and D. Vilone, “Lying on networks: The role of structure and topology in promoting honesty,” *Phys. Rev. E* **101**, 032305 (2020).

⁶⁶A. Kumar, S. Chowdhary, V. Capraro, and M. Perc, “Evolution of honesty in higher-order social networks,” *Phys. Rev. E* **104**, 054308 (2021).

⁶⁷D. Fudenberg and D. K. Levine, *The Theory of Learning in Games* (The MIT Press, Cambridge, 1998).

⁶⁸M. F. Norman, “Some convergence theorems for stochastic learning models with distance diminishing operators,” *J. Math. Psychol.* **5**(1), 61–101 (1968).