

Research



Cite this article: Saponara M, Fernández Domingos E, Pacheco JM, Lenaerts T. 2025 Evolution favours positively biased reasoning in sequential interactions with high future gains. *J. R. Soc. Interface* **22**: 20250153. <https://doi.org/10.1098/rsif.2025.0153>

Received: 21 February 2025

Accepted: 3 July 2025

Subject Category:

Life Sciences—Mathematics interface

Subject Areas:

evolution, biomathematics, computational biology

Keywords:

cognitive bias, level-*k* thinking, evolutionary game theory, sequential dilemmas, mixed-motive games

Author for correspondence:

Marco Saponara
e-mails: marco.saponara@ulb.be;
marcosaponara97@gmail.com

Evolution favours positively biased reasoning in sequential interactions with high future gains

Marco Saponara¹, Elias Fernández Domingos^{1,2}, Jorge M. Pacheco^{3,4} and Tom Lenaerts^{1,2,5}

¹Machine Learning Group, Université Libre de Bruxelles, Campus de la Plaine, Brussels, Belgium

²Artificial Intelligence Lab, Vrije Universiteit Brussel, Brussels, Belgium

³INESC-ID, IST-Tagusparque, Porto Salvo, Portugal

⁴ATP-group, Porto Salvo, Portugal

⁵Center for Human-Compatible AI, UC Berkeley, Berkeley, CA, USA

MS, 0009-0008-8092-6220; EFD, 0000-0002-4717-7984; JMP, 0000-0002-2579-8499; TL, 0000-0003-3645-1455

Empirical evidence shows that human behaviour often deviates from game-theoretical rationality. For instance, humans may hold unrealistic expectations about future outcomes. As the evolutionary roots of such biases remain unclear, we investigate here how reasoning abilities and cognitive biases coevolve using the evolutionary game theory. In our model, individuals in a population deploy a variety of unbiased and biased level-*k* reasoning strategies to anticipate others' behaviour in sequential interactions, represented by the incremental centipede game. Positively biased reasoning strategies have a systematic inference bias towards higher but uncertain rewards, while negatively biased strategies reflect the opposite tendency. We find that selection consistently favours positively biased reasoning, with rational behaviour even going extinct. This bias coevolves with bounded rationality, as the reasoning depth remains limited in the population. Interestingly, positively biased agents may coexist with non-reasoning agents, thus pointing to a novel equilibrium. Longer games further promote positively biased reasoning, as they can lead to higher future rewards. The biased reasoning strategies proposed in this model may reflect cognitive phenomena like wishful thinking and defensive pessimism. This work therefore supports the claim that certain cognitive biases, despite deviating from rational judgement, constitute an adaptive feature to better cope with social dilemmas.

1. Introduction

A large body of experimental evidence has shown that humans can behave differently from the self-interested rational decision-makers hypothesized by game-theoretical models [1–5]. In fact, our decisions may not produce outcomes that align with classical solution concepts such as the Nash equilibrium [6]. Instead, they may contain biases, i.e. consistent deviations from a reference—rational behaviour—deemed to be correct [7–9]. Such biases, on the one hand, influence strategic reasoning and, on the other hand, render their evolutionary roots difficult to explain [10].

Trivers' *theory of self-deception* [11,12] argues that the evolution of cognitive biases stems from the adaptive advantages of deceiving oneself to better persuade or deceive others, thereby reducing the cognitive cost associated with dishonesty in social interactions. This capacity is closely related to the notion of *theory of mind* (ToM) [13], which allows us to attribute mental states such as beliefs, intentions and desires to others [14]. In fact, this early process enabled the development of a crucial feature of human social cognition that is usually defined as *inter-subjectivity* (i.e. an understanding of the intentionality of others),




Figure 1. Extensive form of the six-step incremental centipede game with exponential growth. The notation $\pi_i(t)$ denotes the payoff player i would get if the game ends at step t .

thus rendering social transgressions more likely to be detected and punished.

In a recent publication [15], we proposed an evolutionary model to study the conditions under which ToM may have evolved, finding that intermediate levels of ToM, interpreted as recursive reasoning of the sort *A thinks that B thinks that A thinks that B will do X* [13,16], coevolve with an optimism bias regarding the expected behaviour of other individuals in the context of mixed-motive interactions. This theoretical work also showed that the expectation of higher future rewards is essential for the appearance of ToM, and that the behaviour generated by the surviving reasoning strategies aligns well with previous experimental data [2].

Here, we go beyond the work in [15] by exploring the evolutionary dynamics of biased and unbiased reasoning strategies using the framework of evolutionary game theory [17–20]. Strategic reasoning is represented, as in [15], through the lens of level- k theory [21–23], a popular approach to model *bounded rationality* [24] (i.e. the idea that people have limited cognitive capacity and might follow a strategy that is not perfectly rational). In this framework, the degree of rationality, or *sophistication*, of an individual is represented by an integer $k \geq 0$, where no reasoning ($k = 0$) corresponds to an individual acting on ingrained beliefs, while higher values translate into a more sophisticated reasoning capacity.

The goal here is to investigate whether reasoning strategies with particular properties may be favoured by evolution. In our evolutionary model, described in §4, individuals are equipped with reasoning strategies whose cognitive bias within the reasoning process is modulated by a noise parameter ϵ . Depending on the type of reasoning process, the parameter ϵ perturbs the inferred action away from the deterministic best response at each reasoning step (see §4 for details). As a result, the outcome of the reasoning process of an agent may systematically deviate from what is expected from their reasoning level- k , basing their strategic choices on a subjective interpretation of reality, possibly triggered by both external and internal factors such as the temptation for higher future rewards, as well as the anticipated behaviour of other individuals in the population.

We therefore relax the underlying assumption in [15], wherein every individual shared the same reasoning process and made mistakes in the same way. In this regard, our aim is to answer the following series of questions:

- under which conditions rational behaviour is replaced by different reasoning strategies;
- whether cognitively biased forms of reasoning are preferred to unbiased forms (with biases representing different types of systematic mistakes within each reasoning process); and
- whether an evolutionary robust strategy emerges or whether multiple reasoning strategies may coexist in the population.

As in [15], the current analysis is also performed within the context of the centipede game [25], a sequential game with complete and perfect information that involves two individuals¹ who take turns, over a total of L steps, deciding between two actions ($T = \text{take}$, $P = \text{pass}$) regarding the split of a resource. Player 1 (player 2) can play T at even (odd) steps. Playing T at a given step means ending the game and receiving a larger share of the resource than the co-player. Playing P means letting the other player decide what to do in the next step, unless the last decision node is reached, where player 2 has to decide between two different splits (see figure 1 for an example).

The payoff structure of a centipede game is designed such that solving the game via the game-theoretical method of backward induction leads a rational self-interested individual to play T at every opportunity. Hence, in the unique subgame perfect Nash equilibrium (SPE), each player chooses to end the game as early as possible.² Nevertheless, several studies have shown that this equilibrium strategy is rarely observed during behavioural experiments, especially when the resource to be shared grows with each step of the game [2,28,29]. Here, we focus on the *incremental centipede game* (ICG), where the resource starts at 0.5 and doubles at every step [2,30]. In figure 1, we visualize the extensive form and the payoff structure for the case where $L = 6$.

Several explanations have been proposed to bridge the gap between theoretical and experimental results in the centipede game. Typically, they can be categorized into two main groups, focusing either on insufficient cognitive ability to perform backward induction reasoning [21,22,31,32], or on other-regarding motives that might interfere with self-interested decision-making [33–35]. Here, we argue that the potential evolutionary advantage of certain cognitive biases might have interfered with the development of backward induction reasoning, as these deviations might have been the result of heuristics adapted to deal with the complexity of uncertain environments [36,37].

Our contributions are threefold. First, we show that a reasoning strategy with a systematic inference bias towards higher but uncertain rewards is favoured and becomes dominant under strong selection, whereas rational behaviour undergoes extinction. Individuals employing such a biased inference process think, at each step, that higher payoff outcomes are more likely to occur compared with individuals using the analogous unbiased strategy. In this sense, their reasoning process, if compared with the

¹For extensions to three-player centipede games, see [26,27].

²Remarkably, all Nash equilibria in a centipede game require player 1 to end the game at the first step.

Table 1. Representation of backward induction in a six-step centipede game as a level- k recursive reasoning process. From left to right, each column contains the actions of player 1 (top) and player 2 (bottom) as they are determined by a deterministic introspective process with an increasing number of reasoning steps. Such a process starts at $k = 0$ (left), where players do not perform any reasoning and will simply end the game in accordance with the highest personal payoff. For $k > 0$, players choose the best response to a fictional level—($k - 1$) player. When $k = 5$ (right), players perform the full backward induction process; thus, the action generated by the reasoning process corresponds to the outcome predicted by the subgame perfect Nash equilibrium (SPE).

	$k = 0$	$k = 1$	$k = 2$	$k = 3$	$k = 4$	$k = 5$
Pl.1	$\operatorname{argmax}_{0 \leq t \leq 6} \pi_1(t) = 6$	$\mathcal{BR}_1(5) = 4$	$\mathcal{BR}_1(5) = 4$	$\mathcal{BR}_1(3) = 2$	$\mathcal{BR}_1(3) = 2$	$\mathcal{BR}_1(1) = 0$
Pl.2	$\operatorname{argmax}_{0 \leq t \leq 6} \pi_2(t) = 5$	$\mathcal{BR}_2(6) = 5$	$\mathcal{BR}_2(4) = 3$	$\mathcal{BR}_2(4) = 3$	$\mathcal{BR}_2(2) = 1$	$\mathcal{BR}_2(2) = 1$

unbiased reasoning type, systematically overestimates the probability that the ICG will end at a later node. As a result, they tend (with some probability that depends on the noise parameter ϵ) to direct their decisions to stop the game at later steps with respect to an unbiased level- k process, with no guarantee of actually reaching them. Given these features, we label this reasoning strategy as a *positively biased* reasoning strategy, whereas its inverse (i.e. directing their decisions to stop the ICG at earlier steps with respect to the unbiased reasoning type) is a *negatively biased* reasoning strategy.

Our second contribution is to show that the sophistication of each individual, i.e. the depth of their reasoning processes, remains limited, underlining once more that explicit costs are not required to observe the emergence of bounded rationality [15,38,39]. Interestingly, we reveal the possibility of a coexistence between a myopic payoff-maximizing strategy (i.e. ending the game in accordance with the personal highest payoff without any reasoning about the possible decision of the other player) and one level of positively biased reasoning (i.e. an individual with $k = 1$ reasoning capacity whose mistakes in reasoning are biased towards higher future gains). This finding is aligned with previous econometric analyses [2,22], which assumed the presence of a minority of altruistic players to align with experimental data. Yet, we propose here that some individuals are driven by a desire to obtain the maximum payoff in the game rather than by a high pro-social orientation, thereby not relying on explanations related to other-regarding motives.

Finally, we identify that the length of the ICG, defined by the parameter L , is instrumental in the evolutionary success of positively biased reasoning, as longer exchanges lead to exponentially higher rewards. Specifically, higher values of L amplify the range of cognitive noise ϵ for which positively biased reasoning is the most frequent strategy in the population. We observe, nevertheless, a limit to the window of noise values wherein this reasoning strategy is successful, as excessive cognitive noise renders any type of reasoning evolutionary detrimental.

Through these contributions, the present work further explains the emergence of the non-rational behaviour that is frequently observed in the experimental literature involving this sequential dilemma [34] as well as others [40,41]. Our theoretical model, applied in the context of reciprocal exchanges of resources, thus corroborates the idea that certain cognitive biases, despite leading to systematic deviations from a rational judgement, can constitute an adaptive feature to successfully interact with our peers in this kind of social dilemma.

2. Results and discussion

To understand which form of reasoning, either unbiased or biased, is preferred by evolution, we examine the evolutionary dynamics of five types of reasoning. We consider two basic types, namely the rational strategy of stopping the game as early as possible (i.e. the sub-game perfect equilibrium strategy, or SPE) and the myopic payoff-maximizers who ignore the possible choices of the other player and stop the game at the node that maximizes their personal reward (i.e. no-reasoning, NR , or equivalently $k = 0$). In our case, the latter strategy corresponds to always playing *pass* for player 1 and playing *take* at the last decision node for player 2 (see table 1). Additionally, we consider an unbiased (level- k) reasoning type that recursively applies (k times) a best response perturbed by uniform cognitive noise over the player's action space, along with two biased (level- k) reasoning types. The biased reasoning types are implemented similarly to the unbiased one, but the cognitive noise now directs the individual to end the game later or earlier than would be obtained by noiseless level- k reasoning in the ICG. As explained in §1 the myopic payoff-maximizers, the former will be referred to as a *positively biased* reasoning type, since each step of the reasoning process is biased by a desire for future outcomes with a higher payoff (corresponding to later steps in the ICG) compared with what would be calculated by noiseless level- k reasoning. In opposition, this makes the latter a *negatively biased* reasoning type, as it is preferred to take before the step calculated by noiseless level- k reasoning. Details are provided in §4.

2.1. Positively biased reasoning coevolves with bounded rationality

Figure 2 answers our first question: we show the outcome of the coevolution between reasoning types and k -levels as dictated by the stationary distribution of the stochastic evolutionary model for different values of the selection strength β (see §4; [15,30]). To calibrate the output of our model, we determine the relevant level of cognitive noise ϵ , comparing our predictions with the experimental data from Kawagoe & Takizawa [22]. The best fitting is provided for $\epsilon \approx 0.19$, as shown in figure 5 of appendix A. For comparison purposes, the same analysis and fitting are performed with the data on the same game from the seminal work by McKelvey & Palfrey [2], producing $\epsilon \approx 0.06$ (see appendix A for further details).

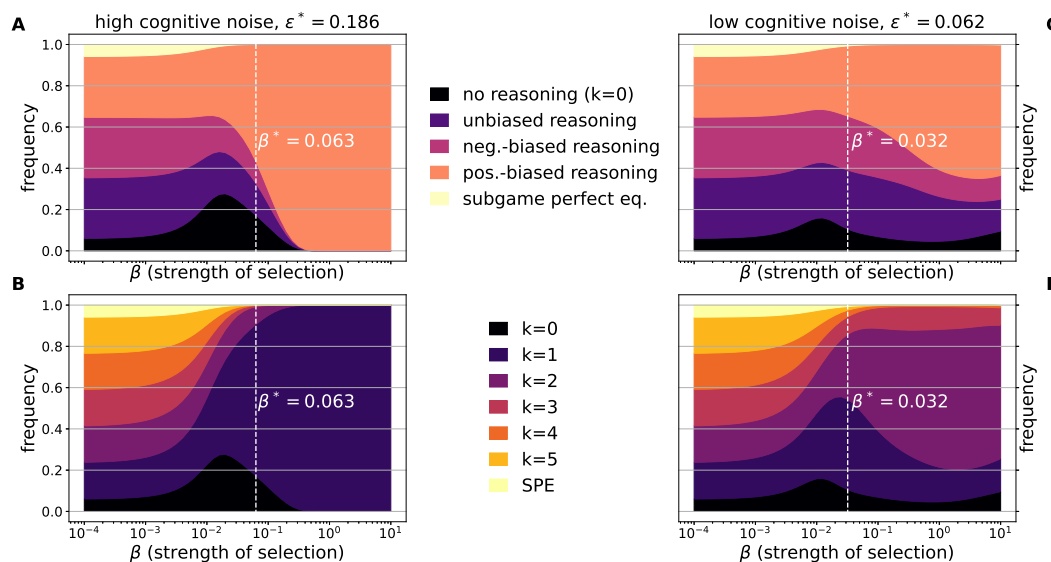


Figure 2. Positively biased reasoning coevolves with bounded rationality. The distribution of reasoning types and reasoning levels in the six-step ICG is shown for variable selection strength β , under high ($\epsilon^* = 0.186$) and low ($\epsilon^* = 0.062$) probability of reasoning errors (panels A,B and C,D, respectively). For each value of β , we compute the stationary distribution ϕ of our dynamical system under the assumption of rare mutations. The frequencies of each strategy are aggregated by reasoning kernel, $f_{\text{ker}} = \sum_{k=1}^L \phi_{s_{\text{ker}},k}$ for $\text{ker} \in \{U, B^+, B^-\}$, and by reasoning level, $f_k = \sum_{\text{ker} \in \{U, B^+, B^-\}} \phi_{s_{\text{ker}},k}$ for $1 \leq k \leq L$ (panels A,C and B,D, respectively). For completeness, the figures also include the frequency of the pure strategy associated with no reasoning (i.e. reasoning level $k = 0$) and the SPE strategy of stopping the game as early as possible. The white dotted lines correspond to the selection strength β^* of optimal fitting with the data in [22] and [2] (panels A,B and C,D, respectively). The population size is fixed to 100 individuals.

In figure 2, panel A shows the abundance of reasoning types as given by the stationary distribution for different values of selection strength β and the calibrated cognitive noise $\epsilon \approx 0.19$. The aim here is to understand how increasing the selective pressure (i.e. higher β) affects the likelihood of ending up with any of the reasoning strategies studied here. When $\beta \approx 0$, the evolutionary dynamics corresponds to the process of neutral drift, thus each reasoning strategy is almost equally present in the population. As β grows, the payoff differences between strategies become increasingly relevant. As can be observed, the positively biased reasoning type takes over the population, fully dominating when $\beta \geq 0.4$. This preference for positively biased reasoning grows together with a preference for performing $k = 1$ reasoning steps, as can be seen in panel B. In fact, after a minor peak in the frequency of $k = 0$ individuals, the $k = 1$ type eventually takes over the entire population.

The right panels in figure 2 show the results for the same model fitted to the experimental data in [2]. With a lower value of ϵ (i.e. $\epsilon \approx 0.06$), the mistakes within the reasoning processes are less frequent than in the previous case. The differences between reasoning types are therefore not as marked. However, positively biased reasoning is again the most frequent type under strong selection ($\beta \geq 1$), being adopted by more than 60% of the population (panel C). Moreover, the increased level of reasoning accuracy allows for the development of a higher reasoning capacity, with $k = 2$ now being the most frequent level of reasoning, adopted by 60 to 68% of the population under strong selection, followed by $k = 1$ and $k = 3$, both with comparable frequencies (panel D).

In §2.2, we examine the stochastic evolutionary dynamics that leads to this result in more detail. Before, we further observe that as the stochasticity of the imitation process decreases (i.e. for larger values of β), higher levels of reasoning ($k = 4$ and $k = 5$) progressively disappear from the population, meaning that the sophisticated individuals are replaced by the less sophisticated ones. This result is in agreement with both theoretical [15,38] and experimental [16,22] findings. Notably, higher levels of reasoning do not need to be explicitly penalized through the introduction of any computational cost because every additional backward step in the reasoning process has the implicit cost of propagating and amplifying the effect of cognitive noise. Therefore, the population eventually finds the best trade-off between sophistication and propagation of errors, leading to the presence of intermediate reasoning levels $k = 1, 2, 3$ when ϵ is sufficiently small and to the invasion of the level $k = 1$ at higher ϵ .

Remarkably, the pure strategy associated with the SPE (i.e. to stop the game as early as possible) never manages to invade the others. This novel outcome is at odds with a previous evolutionary model proposed by Rand & Nowak [30], which showed that this strategy would eventually prevail under strong selection when the population only plays pure strategies. Our model avoids this outcome because each strategy is a noisy introspective process [41,42], where iterated conjectures about others' decisions might undergo stochastic deviations from the deterministic rational predictions. This result therefore underlines once again the importance of the introduction of cognitive noise, and therefore of the generated mixed strategies, to obtain a more flexible and insightful model.

2.2. A coexistence dynamics emerges from biased reasoning

In figure 3, we perform a more detailed analysis of the evolutionary dynamics to explain the previous observations. Panel A shows the invasion diagram (i.e. the reduced Markov chain connecting all possible monomorphic population states) for the calibrated

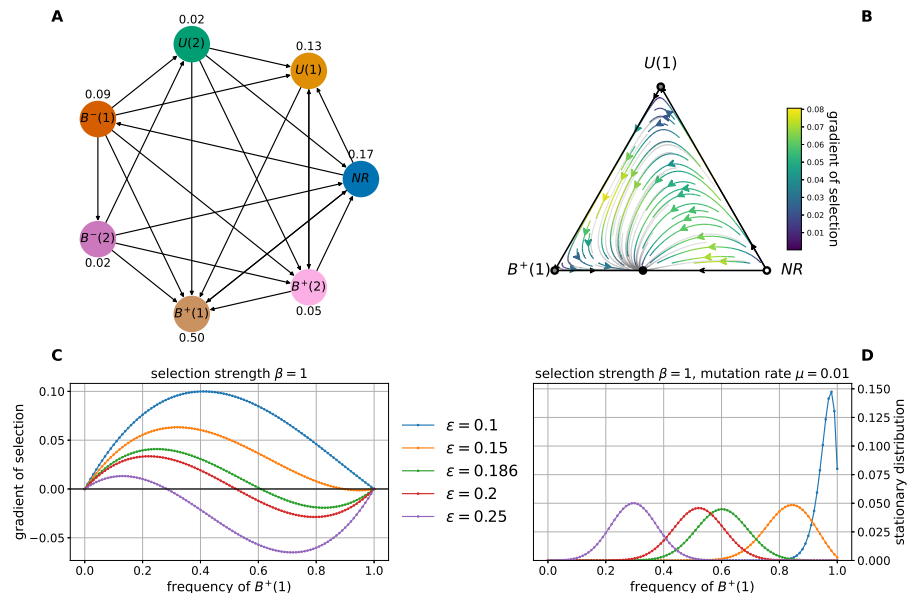


Figure 3. No-reasoning and level-1 positively biased reasoning coexist under limited cognitive noise. Panel A shows the invasion diagram related to the evolutionary dynamics of our model under the small mutation limit, for $\beta = 0.063$ and $\varepsilon = 0.186$. The strategies are limited to $k \leq 2$, and when $k > 0$ the reasoning kernel can be unbiased (U), positively biased (B^+) or negatively biased (B^-). The edges in the graph represent the transitions between different monomorphic states: an outgoing edge from strategy A to B means that A is invaded by B . The number next to each node is the fraction of time spent in each monomorphic state at the stationary regime. Panel B focuses on the evolutionary dynamics in an infinite population between the three most frequent strategies in panel A, i.e. $B^+(1)$, NR and $U(1)$. We plot the gradient of selection and the fixed points associated with the replicator equation. The black and white circles represent stable and unstable equilibria respectively, whereas grey circles are saddle points. The arrows indicate the direction of the selective pressure. Finally, panels C and D show the effect of the cognitive noise ε on the location of the coexistence point between the two strategies $B^+(1)$ and NR under strong selection ($\beta = 1$). The population size in panels A, C and D is $Z = 100$.

values $\varepsilon = 0.19$ and $\beta = 0.063$, corresponding to the settings indicated with the white dotted lines in figure 2A,B. For a better visualization, the strategies are restricted to $k \leq 2$, as higher k -levels cover a negligible portion of the stationary distribution. The invasion diagram is a useful tool to understand the dominance relationships between strategies under the small-mutation limit approximation. A directed edge from node i to node j in this diagram indicates that a single mutant adopting strategy j is able to fixate in a population of i s with a probability higher than random drift [19]. An evolutionary robust strategy (ERS) is a strategy with no outgoing arrows [43].

As we can see, for these parameter settings, there does not exist any ERS.³ Instead, an interesting cycle is present, involving the three most frequent strategies, namely the positively biased reasoning strategy with $k = 1$, $B^+(1)$, the no-reasoning (NR) or payoff-maximizing strategy and $k = 1$ level of unbiased reasoning $U(1)$. What is interesting is that the arrow connecting $B^+(1)$ and NR is pointing in both directions, which may indicate a potential for coexistence of both reasoning strategies [44].

This hypothesis is confirmed in panel B, where we consider the replicator dynamics [45,46] restricted to the three strategies NR , $B^+(1)$ and $U(1)$, in which the population spends approximately 80% of the time according to the stationary distribution visualized in panel A. The results of the replicator dynamics are visualized through the simplex covering all population configurations for the three aforementioned strategies, including the anticipated dynamics in each point of the simplex. One can clearly observe that a stable coexistence point appears for the calibrated cognitive noise $\varepsilon \approx 0.19$, wherein 60% of the population uses level-1 positively biased reasoning ($B^+(1)$), while the remaining 40% uses NR . At each interior point of the simplex, the replicator dynamics is directed towards this novel equilibrium. We thus see that the game and its outcomes are transformed due to the presence of biased and unbiased reasoning strategies [15,47,48], revealing the impact that (biased) reasoning strategies may have on decision-making.

This transformation of the game and resulting coexistence dynamics is further explored in panels C and D of figure 3, where the complete Markov chain connecting the monomorphic population with NR to the one containing only $B^+(1)$ is analysed. Both plots show the interplay between the intensity of cognitive noise ε and the coexistence point when the game is restricted to the two aforementioned strategies. As we can see, when ε is sufficiently small ($\varepsilon \approx 0.1$), strong selection leads to the full invasion of individuals adopting the strategy $B^+(1)$. Larger values of ε allow the emergence of a coexistence point, where the fractions of the NR type increases with ε , until the amount of cognitive errors is so large that performing any reasoning becomes evolutionary detrimental. The results remain robust when studying the full Markov chain under mutations (see panel D).

Overall, these results confirm the importance of the positively biased reasoning strategy within the context of the ICG. Indeed, the possibility of developing positively biased reasoning allows the population to enter a new stable equilibrium, where the majority of the population is made of positively biased reasoners that coexist with a minority of myopic payoff-maximizers (assuming a sufficiently small probability of reasoning mistakes). This novel equilibrium, while leading to a better agreement with the experimental reference, is also aligned with previous theoretical and behavioural studies arguing that humans' sophistication in a population is often heterogeneous and that 'naive' types can survive and coexist with more sophisticated individuals [16,49,50].

³As mentioned previously, when $\varepsilon \approx 0.19$ the positively biased reasoning type with $k = 1$ turns into an ERS for $\beta > 0.4$.

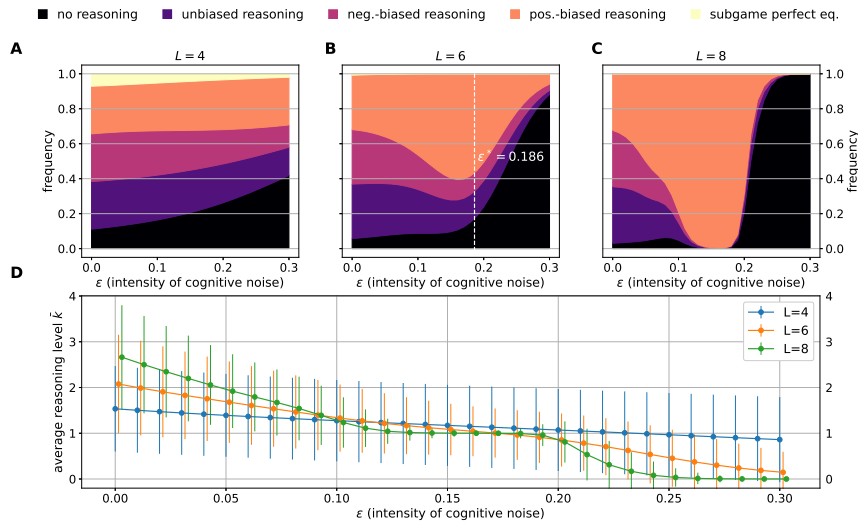


Figure 4. Longer exchanges promote the emergence of positively biased reasoning. The effect of the number of steps of the ICG, represented by the parameter $L \in \{4, 6, 8\}$, on the evolution of reasoning kernels (panels A–C) and reasoning levels (panel D) is shown for variable cognitive noise ϵ . For each value of ϵ , we compute the stationary distribution ϕ of our dynamical system under the assumption of rare mutations. In panels A–C, the frequencies of each strategy are aggregated by reasoning kernel, $f_{\text{ker}} = \sum_{k=1}^L \phi_{\text{ker},k}$ for $\text{ker} \in \{U, B^+, B^-\}$. For completeness, the figures also include the frequency of the pure strategy associated with no reasoning (i.e. reasoning level $k = 0$) and the SPE strategy of stopping the game as early as possible. In panel D, we report the average reasoning level of the population, $\bar{k} = \sum_{k=1}^L k f_k$, where $f_k = \sum_{\text{ker} \in \{U, B^+, B^-\}} \phi_{\text{ker},k}$, along with its standard deviation plotted in the form of error bars. The white dotted line in panel B corresponds to the value of cognitive noise ϵ^* of optimal fitting with the data in [22]. The population size is fixed to 100 individuals, and the strength of selection is $\beta^* = 0.063$.

2.3. Longer exchanges promote the emergence of positively biased reasoning

The length of the ICG, represented here by the parameter L , is a key factor to secure the evolutionary success of positively biased reasoning. Indeed, longer ICGs translate into a much larger resource to be shared among the players because of the exponential growth inherent in the payoff structure. In figure 4, we explore the effect of the intensity of cognitive noise ϵ on the stationary distribution of our model when the selection strength β is fixed at 0.063, for three values of L ($L = 4, 6, 8$ in panels A, B, C, respectively).

When $L = 4$, the difference in frequency between the unbiased, positively biased and negatively biased reasoning is negligible, regardless of the intensity of the cognitive noise. When $L = 6$, the incentive to end the game at later nodes is exponentially higher, leading to a prevalence of positively biased reasoning for a window of ϵ values: in fact, more than half of the population develops a positivity bias within the range $0.11 < \epsilon < 0.2$. Finally, when $L = 8$, the same effect is even more pronounced: now, the region where more than half of the population chooses positively biased reasoning extends to the range $0.05 < \epsilon < 0.2$, and in the subset of values $0.14 < \epsilon < 0.17$ the entire population adopts the strategy previously mentioned $B^+(1)$.

In each of the three scenarios, however, a larger probability of cognitive errors renders the reasoning process increasingly harmful, regardless of the reasoning process employed. This is particularly evident in the plot showing the decay of the average reasoning level \bar{k} (figure 4D). The figure reports the mean and the standard deviation of the distribution of reasoning levels within the population in the stationary state as a function of cognitive noise. We can see how, as ϵ increases, the myopic payoff-maximizers progressively spread across the population. This effect is more pronounced as the length of the ICG increases to $L = 8$.

Once again, the presence of the strategy associated with the SPE is negligible, regardless of the value of ϵ . This result is remarkable because this pure strategy, despite being deterministic, is indirectly affected by the amount of mistakes committed within the population of interacting agents. In a population of sophisticated individuals who, nonetheless, make reasoning mistakes, it is better to take advantage of these mistakes and opt for a myopic maximization of one's own payoff, rather than performing the rational reasoning process of backward induction.

3. Conclusion

We presented a model based on the ICG to study the coevolution of strategic reasoning with and without two cognitive biases, which were referred to as positively and negatively biased reasoning, in the framework of evolutionary game theory. We showed how evolutionary dynamics in finite populations can lead boundedly rational players in this game to develop a positivity bias while recursively reasoning about the potential move of their opponents. Furthermore, the emergence of positively biased reasoning in this sequential setting appears to be strongly linked to two factors, namely the presence of individuals acting as myopic payoff-maximizers, and the promise of higher future gains (captured here by the length of the ICG and the exponential growth of the resource to be shared among the players).

While our work was restricted to an evolutionary analysis of positively and negatively biased reasoning strategies in the ICG, the model may be interpreted to reflect specific cognitive and psychological phenomena. As an example, the aforementioned

reasoning strategies could be potentially associated with *wishful thinking* [51–53] and its counterpart, *defensive pessimism* [54–56]. Wishful thinking is described in the literature as a distortion of cognitive forecasting processes towards more optimistic expectations, and it implies the overestimation of the likelihood of desirable events. This concept therefore fits, to a certain extent, our model, as the ICG formalizes the partition of a growing resource between parties that hold beliefs about each other and have desires for possibly different outcomes.

Wishful thinking, however, is an extremely broad phenomenon. Moreover, the related experimental literature exploring the link between wishful thinking and strategic decision-making is rather limited and produced mixed results (e.g. [57,58] for a review). It is therefore not obvious that wishful thinking is the source of the behavioural deviations presented here or that the model truly captures the intricacies of this concept. Alternative interpretations, such as a risk-taking tendency [59,60], may also align with the results of our model. A more comprehensive analysis of the link between the model and these cognitive interpretations is therefore needed, which is left for future work.

Our study was also limited to the investigation of cognitive biases in the form of consistent deviations from a rational judgement of the opponent's future choice. The cognitive bias of each player was unequivocally translated into a specific behaviour. In this sense, our study assumed that the relationship between behavioural strategies and cognitive biases is unambiguous, while each strategy could, in reality, be the outcome of multiple cognitive processes and heuristics [36], as remarked above. Moreover, our model assumed a static belief framework in which individuals hold fixed prior beliefs about their co-players. It also imposes that individuals share the same beliefs throughout the population and that all players use the same type of recursive reasoning process, where opponents always possess a comparatively lower capacity for reasoning [15]. Future work can relax these assumptions to examine how robust the results remain while also opening the door to alternative models to explore other aspects of cognitive biases and related psychological phenomena that were not yet considered here.

To conclude, the insights provided by our model, while confirming previous experimental [22,61] and theoretical [15,38,52] results, show that the deviations from the rational behaviour prescribed by game theory can be the outcome of evolutionary processes in which reasoning heuristics were developed in order to better interact in complex and uncertain environments. In particular, positively biased reasoning and bounded rationality proved to be adaptive features in the context of reciprocal exchanges of a shared resource that grows over time, in line with theories of the evolution of human social cognition claiming that we developed these non-rational behaviours by means of our remarkable ability to understand the intentionality of our peers.

4. Methods

4.1. Notation

Let us first introduce the notation that we will use throughout this section:

- the subscript $i \in \{1, 2\}$ denotes the role of the player, while the classical notation $-i$ denotes the role of the co-player;
- \mathcal{A}_i represents the set of possible actions of the player in role i : specifically, $\mathcal{A}_1 = \{0, 2, \dots, L-2, L\}$ and $\mathcal{A}_2 = \{1, 3, \dots, L-1, L\}$, with $|\mathcal{A}_1| = |\mathcal{A}_2| = 1 + L/2$ (note that the action L , shared by both players, means to always play *pass*);
- $T_i \in \mathcal{A}_i$ denotes the step at which the player in role i will end the game; and
- σ_i represents the probability distribution over the different steps at which the player in role i may end the game.

4.2. Strategies

Given an ICG of L steps, our game-theoretical model considers two pure strategies and $3 \times (L-1)$ generative strategies.⁴ The pure strategies consist of the sub-game perfect Nash equilibrium (SPE) strategy and a myopic payoff-maximization strategy. The former will decide to end the game as early as possible (in line with full backward induction). The latter aims at acquiring the highest possible personal reward, ignoring the potential choices of the co-player. If $L = 6$, for instance, when acting as player 1, this strategy will play $T_1 = 6$ because $\pi_1(6) = 25.6$ is the highest payoff of the game for player 1 (see figure 1), without taking into account the other player's likely *take* move at the fifth node. Likewise, a myopic player 2 would stop the game at the step $T_2 = 5$ to obtain the maximum payoff $\pi_2(5) = 12.8$, without considering that player 1 would probably stop the game at the preceding node (see figure 1 and table 1).

The $3 \times (L-1)$ generative strategies are defined using the level- k framework [15,21]. Each generative strategy is determined by a value for k and the type of (un)biased reasoning they will use to arrive at a decision. Each reasoning level k corresponds to k steps of backward induction reasoning, where players might deviate from the rational best response at each reasoning step due to cognitive noise, defined by a parameter $\varepsilon \in [0, 1]$. At $k = 0$, which corresponds to no-reasoning (NR), the strategy corresponds to the myopic payoff-maximizer explained earlier. When $k > 0$, level- k individuals perform a recursive reasoning process, assuming that the co-player belongs to the type $k-1$.

In a situation, where individuals do not make any reasoning mistakes ($\varepsilon = 0$), this process can be formulated as follows:

$$T_i(k) = \begin{cases} \operatorname{argmax}_{0 \leq t \leq L} \pi_i(t), & \text{if } k = 0 \\ \mathcal{BR}_i(T_{-i}(k-1)), & \text{otherwise} \end{cases}, \quad (4.1)$$

⁴We use the term *generative* as opposed to *mixed* because in each interaction, these strategies need to generate the action based on a noisy level- k recursive process. While they implicitly encode a mixed strategy, i.e. a probability distribution over the action space of both players, they are defined by the reasoning level $k > 0$ and one of the three reasoning types, as it is further explained in §4.

where $\mathcal{BR}_i(\cdot)$ encodes the best response of player i given the action of player $-i$. This case is essentially analogous to [30], with the reasoning level $k = L - 1$ being equivalent to the SPE strategy (see table 1).

Since we are interested in the evolution of cognitive biases and strategic reasoning, we focus on the case $\varepsilon > 0$: in particular, at each reasoning step, we let players deviate from the best response with probability $\varepsilon \in [0, 1]$ [15], i.e. player i will now choose their action according to a noisy best response, $\mathcal{NBR}_i(\cdot)$, which replaces $\mathcal{BR}_i(\cdot)$ in equation (4.1). Resuming the example mentioned above, if player 2 now performs one reasoning step ($k = 1$), they will take into account that a *fictional* player 1 with $k = 0$ follows the pure strategy $T_1(0) = 6$, so player 2 will adopt strategy $T_2(1) = 5$ with probability $1 - \varepsilon$, but will deviate from it with probability ε . Depending on the reasoning type, these deviations can occur uniformly on the whole action space if the reasoning process is unbiased, or they can be skewed towards earlier or later steps in the case of negatively or positively biased reasoning, respectively.

Each reasoning process will thus produce a role-dependent probability distribution over the actions of each player, $\sigma_1(k)$ and $\sigma_2(k)$, corresponding to the probability distributions over the action spaces of player 1 and player 2, respectively,

$$\begin{aligned}\sigma_1(k) &= [\mathbb{P}(T_1(k) = 0), \mathbb{P}(T_1(k) = 2), \dots, \mathbb{P}(T_1(k) = L)], \\ \sigma_2(k) &= [\mathbb{P}(T_2(k) = 1), \mathbb{P}(T_2(k) = 3), \dots, \mathbb{P}(T_2(k) = L)].\end{aligned}$$

The recursive reasoning in equation (4.1) then corresponds to an iterative application of the law of total probability, where each element of the vectors $\sigma_1(k)$ and $\sigma_2(k)$ equals

$$\mathbb{P}(T_i(k) = t_i) = \sum_{t_{-i} \in \mathcal{A}_{-i}} \mathbb{P}(\mathcal{NBR}_i(t_{-i}) = t_i \mid T_{-i}(k-1) = t_{-i}) \cdot \mathbb{P}(T_{-i}(k-1) = t_{-i}), \quad (4.2)$$

with $i \in \{1, 2\}$, $k > 0$ and $t_i \in \mathcal{A}_i$. The overall probability of player i playing action t_i is thus the sum of the probabilities that t_i is chosen as the response to a fictional level- $(k-1)$ opponent playing the action t_{-i} , over all possible actions in the set \mathcal{A}_{-i} .

It can be shown that the strategy of a player covering both roles with equal probability can be computed through the following formula (see appendix B for further details):

$$\sigma(k) = \sigma(0) \cdot M(\varepsilon)^k, \quad (4.3)$$

where $\sigma(k)$ with $0 \leq k \leq L-1$ is the concatenation of the two role-dependent strategy vectors $\sigma_1(k)$ and $\sigma_2(k)$, and $M(\varepsilon)$ is a matrix that depends on the intensity of the cognitive noise ε and encapsulates the noisy best responses for both roles.

Each strategy in our model is therefore uniquely determined by three components of the reasoning process: (i) the starting point, $\sigma(0)$; (ii) the reasoning kernel, $M(\varepsilon)$; and (iii) the depth of the reasoning process, $0 \leq k < L$. This framework has an intuitive and straightforward interpretation: each individual has a *prior behaviour*, $\sigma(0)$, about the action to choose without any reasoning, and they modify it by applying the reasoning kernel M (representative of their *mind*) iteratively for k times. As stated previously, we fix the prior belief $\sigma(0)$ of each individual such that $k = 0$ corresponds to an agent who myopically aims at maximizing the payoff of the game, leading to $\sigma(0) = [0, 0, 0, 1, 0, 0, 1, 0]$ for $L = 6$.

To study the evolution of cognitive biases, we propose the following reasoning kernels:

- an *unbiased* reasoning kernel, $M_U(\varepsilon)$, where the best response is chosen at each step with probability $1 - \varepsilon$ while any other action is chosen with uniform probability $2\varepsilon/L$;
- a *negatively biased* reasoning kernel, $M_{B^-}(\varepsilon)$, where deviations from the best response only happen to shift towards earlier nodes of the game, i.e. lower-payoff outcomes;
- a *positively biased* reasoning kernel, $M_{B^+}(\varepsilon)$, where deviations from the best response only happen in the direction of later nodes, i.e. outcomes where payoffs are higher but uncertain.

We denote the level- k strategies of unbiased, positively biased and negatively biased reasoning by $U(k)$, $B^+(k)$ and $B^-(k)$, respectively. In appendix B, we report the three matrices $M_U(\varepsilon)$, $M_{B^+}(\varepsilon)$ and $M_{B^-}(\varepsilon)$ in the case where $L = 6$.

4.3. Payoffs

We now describe the computation of the payoffs resulting from the interaction between two strategies. Let us suppose two players, Alice and Bob, adopt strategies σ_A and σ_B , respectively. As stated previously, each of these vectors consists of a pair of role-dependent probability distributions, namely $(\sigma_{1A}, \sigma_{2A})$ and $(\sigma_{1B}, \sigma_{2B})$.

Here, we assume that Alice and Bob play against each other, covering both roles with equal probability, so that the game becomes symmetric [30]. This means that, for instance, the expected payoff of Alice, $\Pi_A(\sigma_A, \sigma_B)$, is given by the average between the expected payoff she would obtain while playing as player 1 against Bob playing as player 2 and the expected payoff in the opposite scenario. Thus, we obtain the following equation:

$$\Pi_A(\sigma_A, \sigma_B) = \frac{1}{2} \sum_{t=0}^L \left\{ \pi_1(t) \cdot \mathbb{P}(T_1^{\min} = t) + \pi_2(t) \cdot \mathbb{P}(T_2^{\min} = t) \right\}, \quad (4.4)$$

where T_1^{\min} and T_2^{\min} are the two random variables associated with the distribution of the terminal node of the game on the two admissible scenarios, where Alice (Bob) plays as player 1 (2) and player 2 (1), respectively:

$$T_1^{\min} := \min\{\sigma_{1A}, \sigma_{2B}\}, \quad T_2^{\min} := \min\{\sigma_{2A}, \sigma_{1B}\}.$$

Note that, since these two random variables are independent and defined over the same support, the cumulative distribution function of the minimum between the two can be computed analytically through the following equation:

$$\mathbb{P}(T_i^{\min} \leq t) = 1 - \mathbb{P}(\sigma_{iA} > t)\mathbb{P}(\sigma_{-iB} > t).$$

4.4. Evolutionary dynamics

As mentioned in §1, we adopt an approach based on the evolutionary game theory to understand how cognitive biases and strategic reasoning might have coevolved in the long term. In brief, this approach approximates the evolution of the behaviour of a population of Z individuals interacting with each other according to their given strategies, via analytical methods or numerical simulations [62]. Borrowing ideas from evolutionary biology, a strategy can propagate within the population if it is associated with higher fitness, i.e. yield a higher expected payoff than the others.

At each generation step, one individual X is randomly chosen to update their strategy by either undergoing a stochastic mutation with probability μ or by imitating a better co-player [63]. In the latter case, another individual Y is randomly sampled from the population to act as a potential role model, and will be imitated with probability $p = 1/(1 + \exp(\beta(\Pi_X - \Pi_Y)))$, which increases with the fitness difference between Y and X [64]. The parameter β represents the intensity of selection, i.e. the strength of imitation: when $\beta = 0$, imitation occurs with 50% chance, corresponding to the process of neutral drift; when β is large, imitation becomes sensitive to the slightest fitness difference, leading to an almost deterministic evolutionary process. The number of states of this birth–death process, however, quickly becomes intractable via analytical methods.

A possible solution is the so-called *small-mutation limit*, i.e. the case where mutations are negligible ($\mu \rightarrow 0$) [65]. This approximation allows us to reduce the size of the Markov chain to the number of strategies of the game. Indeed, with this approximation, the time interval between two mutations is sufficiently large that evolution will lead to the fixation of one strategy in the population before the next mutation leads to the appearance of a new strategy. Thus, at any time, there will be at most two strategies simultaneously present in the population. The rare-mutation limit leads to an embedded Markov chain whose states correspond to the different homogeneous configurations of the population in which everyone plays the same strategy. Most results in this work, except for panels B–D of figure 3, are obtained under the small-mutation limit.

A further approximation, used in panel B of figure 3 is given by the limit of infinite populations. The dynamics in this case is represented by the replicator equations [45,46]. This differential equation expresses a deterministic selection process where the frequency of a type i , x_i , increases if it has higher fitness, Π_i , than the average fitness of the population $\bar{\Pi}$. In formulas, $\dot{x}_i = x_i(\Pi_i(x) - \bar{\Pi}(x))$. When the gradient is 0, i.e. $\dot{x}_i = 0$, we have an equilibrium point. If the gradients in a small neighbourhood of an equilibrium x^* point towards x^* , the equilibrium is said to be evolutionary stable. This definition means that an evolutionary stable state is robust to small changes in the population, i.e. a mutant will not drive the population to a different state. This solution concept, more strict than the well-known Nash equilibrium, allows to characterize complex evolutionary processes through relatively simple differential equations.

Ethics. This work did not require ethical approval from a human subject or animal welfare committee.

Data accessibility. The software implementation of the evolutionary processes is based on EGT Tools⁵ [20], an open-source hybrid C++/Python library that provides both analytical and numerical methods to study game-theoretical problems within the framework of Evolutionary Game Theory. The code to reproduce all the results presented in this work can be found at the following GitHub repository [66].

Declaration of AI use. We have not used AI-assisted technologies in creating this article.

Authors' contributions. M.S.: conceptualization, formal analysis, investigation, methodology, software, visualization, writing—original draft, writing—review and editing; E.F.D.: conceptualization, software, supervision, validation; J.M.P.: conceptualization, supervision, writing—review and editing; T.L.: conceptualization, funding acquisition, investigation, methodology, project administration, supervision, visualization, writing—original draft, writing—review and editing.

All authors gave final approval for publication and agreed to be held accountable for the work performed therein.

Conflict of interest declaration. We declare we have no competing interests.

Funding. The authors gratefully acknowledge the research support of the F.R.S-FNRS (project grant 40007793). T.L. further acknowledges the support of the Service Public de Wallonie Recherche (grant 2010235-ARIAC) by DigitalWallonia4.ai and the Flemish Government through the AI Research Program. E.F.D. is supported by an F.W.O. Senior postdoctoral grant (12A7825N).

Acknowledgements. The authors also thank Axel Abels and the anonymous reviewers for their useful comments and suggestions, which allowed us to improve this article.

Appendix A. Fitting with experimental data

We address the choice of the pair of values for the parameters of our model, namely the strength of selection β and the probability of reasoning error ϵ . In order to retrieve reasonable values, we fitted our model to two different experimental references already mentioned in §2: the seminal study by McKelvey & Palfrey [2] and the more recent work by Kawagoe & Takizawa [22]. Both references provide a benchmark for the behavioural profile in the ICG with $L = 6$ depicted in figure 1. The main difference between these studies lies in the fact that the latter behavioural experiment is a one-shot interaction, while the former consists of 10 iterations with random pairing (in this case, we exclusively refer to the frequencies associated with the first round, to eliminate potential learning effects). Cultural differences may also be at play [67–70], as the experiments in [2] and [22] were conducted in American and Japanese universities, respectively.

⁵<https://egtttools.readthedocs.io/en/latest/>

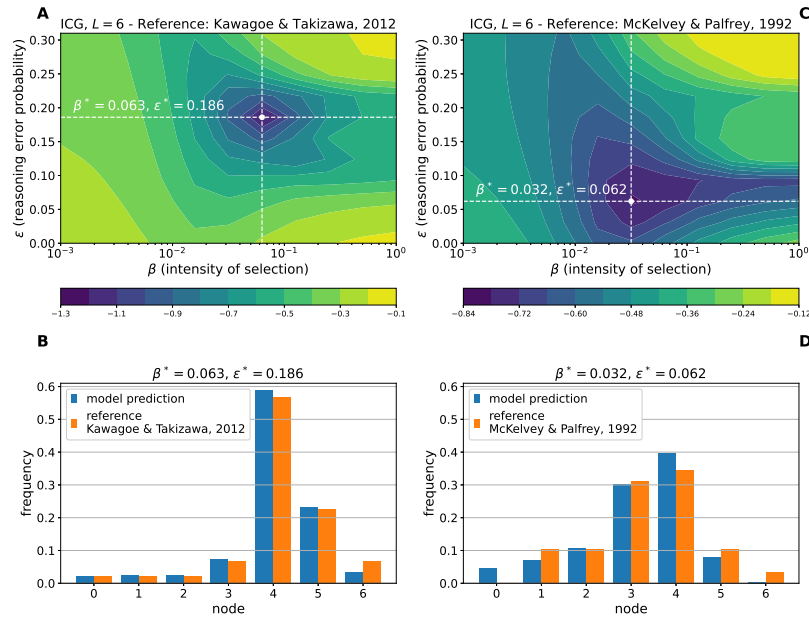


Figure 5. Fitting with experimental references. We fit our model to the experimental data related to the ICG with $L = 6$ from [22] and [2] (panels A,B and C,D, respectively). Panels A,C show the Jensen–Shannon divergence between our model’s prediction and each experimental reference (in logarithmic scale) as a function of the strength of selection β and the probability of reasoning errors ϵ . In the case of [22], the minimum distance (0.06) is achieved when $\beta^* = 0.063$ and $\epsilon^* = 0.186$. When using the data from [2], the minimum distance (0.16) is reached when $\beta^* = 0.032$ and $\epsilon^* = 0.062$. Panels B,D portray the comparison between the optimal fitting and the experimental reference in both cases. The population size is fixed to 100 individuals.

These differences are reflected in the results of the fitting shown in figure 5. By minimizing the Jensen–Shannon divergence between our model’s prediction of the terminal node of the game and the two experimental references, we obtain rather different results. When using the data from [22] (panel A), the best fit is located where the strength of selection and the cognitive noise take values $\beta^* = 0.063$ and $\epsilon^* = 0.186$, respectively. In this case, the relatively high value of ϵ implies that the three different reasoning kernels, and particularly the positively biased one, are necessary to explain the experimental results.

In contrast, when using the data from [2] (panel B), the best fitting is given by lower values, $\beta^* = 0.032$ and $\epsilon^* = 0.062$, respectively. This implies that, in this case, the most relevant factors to explain the data are the stochastic effects of the imitation process (represented by the low strength of selection), as well as the reasoning capacity of each individual (represented by the level k). However, the low value of cognitive noise leads to marginal differences in reasoning types and consequently to almost deterministic strategies. Despite the highly mitigated differences between reasoning kernels, we have seen in figure 2 (panel C) that it is still preferable to develop positively biased reasoning, especially under the effect of stronger selection.

Although *general trust*, intended as an expectation of strangers’ goodwill, has been suggested to be more prevalent in western cultures (particularly Anglophone), Japanese culture has been associated with *assurance-based trust*, which does not require a belief in the benevolence of the other person because it is rather based on the mutual knowledge of the incentive structure surrounding the relationship [67]. This general finding is confirmed by cross-cultural studies involving the centipede game [69,70], where Japanese participants systematically ended the game at nodes later than Western participants. Our results appear to corroborate these previous findings, as approximately 51% of the population in our model adopts the strategy associated with one step of positively biased reasoning when fitted on the Japanese data, while the distribution of strategies is more heterogeneous when the model is fitted on the American data, with one step of positively biased reasoning being nevertheless the most frequent strategy (approximately 18%).

Appendix B. Derivation of the reasoning strategies

Here, we report the mathematical derivation of the strategies of our model. Starting from equation (4.2), it is useful to encode the noisy best responses, $\mathcal{NB}\mathcal{R}_i(\cdot)$, into two role-dependent matrices, $M_1(\epsilon)$ and $M_2(\epsilon)$, which depend on the intensity of the cognitive noise ϵ and contain the probabilities of choosing a certain response (column) given the belief about the action of the co-player (row),

$$[M_i(\epsilon)]_{t_{-i}, t_i} = \mathbb{P}(\mathcal{NB}\mathcal{R}_i(t_{-i}) = t_i \mid T_{-i}(k-1) = t_{-i}).$$

Then, we can rewrite equation (4.2) in the following compact form:

$$\sigma_i(k) = \sigma_{-i}(k-1) \cdot M_i(\epsilon),$$

with $i \in \{1, 2\}$, and $k > 0$. Finally, if we concatenate the two mixed strategy vectors $\sigma_1(k)$ and $\sigma_2(k)$ into one vector $\sigma(k) = [\sigma_1(k), \sigma_2(k)]$ and define a single $(L + 2) \times (L + 2)$ block matrix, $M(\varepsilon)$, such that

$$M(\varepsilon) = \begin{bmatrix} \mathbf{0} & M_2(\varepsilon) \\ M_1(\varepsilon) & \mathbf{0} \end{bmatrix},$$

we obtain the following simple analytical formula encompassing the strategy of a player covering both roles with equal probability:

$$\sigma(k) = \sigma(0) \cdot M(\varepsilon)^k.$$

As mentioned in §4, we consider three reasoning kernels, namely $M_U(\varepsilon)$, $M_{B^-}(\varepsilon)$ and $M_{B^+}(\varepsilon)$ for unbiased, negatively and positively biased reasoning, respectively. Here, we report the three matrices for the case $L = 6$:

$$M_U(\varepsilon) = \begin{bmatrix} 0 & 0 & 0 & 0 & 1/4 & 1/4 & 1/4 & 1/4 \\ 0 & 0 & 0 & 0 & 1-\varepsilon & \varepsilon/3 & \varepsilon/3 & \varepsilon/3 \\ 0 & 0 & 0 & 0 & \varepsilon/3 & 1-\varepsilon & \varepsilon/3 & \varepsilon/3 \\ 0 & 0 & 0 & 0 & \varepsilon/3 & \varepsilon/3 & 1-\varepsilon & \varepsilon/3 \\ 1-\varepsilon & \varepsilon/3 & \varepsilon/3 & \varepsilon/3 & 0 & 0 & 0 & 0 \\ \varepsilon/3 & 1-\varepsilon & \varepsilon/3 & \varepsilon/3 & 0 & 0 & 0 & 0 \\ \varepsilon/3 & \varepsilon/3 & 1-\varepsilon & \varepsilon/3 & 0 & 0 & 0 & 0 \\ \varepsilon/3 & \varepsilon/3 & \varepsilon/3 & 1-\varepsilon & 0 & 0 & 0 & 0 \end{bmatrix},$$

$$M_{B^-}(\varepsilon) = \begin{bmatrix} 0 & 0 & 0 & 0 & 1/4 & 1/4 & 1/4 & 1/4 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \varepsilon & 1-\varepsilon & 0 & 0 \\ 0 & 0 & 0 & 0 & \varepsilon/2 & \varepsilon/2 & 1-\varepsilon & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \varepsilon & 1-\varepsilon & 0 & 0 & 0 & 0 & 0 & 0 \\ \varepsilon/2 & \varepsilon/2 & 1-\varepsilon & 0 & 0 & 0 & 0 & 0 \\ \varepsilon/3 & \varepsilon/3 & \varepsilon/3 & 1-\varepsilon & 0 & 0 & 0 & 0 \end{bmatrix},$$

$$M_{B^+}(\varepsilon) = \begin{bmatrix} 0 & 0 & 0 & 0 & 1/4 & 1/4 & 1/4 & 1/4 \\ 0 & 0 & 0 & 0 & 1-\varepsilon & \varepsilon/3 & \varepsilon/3 & \varepsilon/3 \\ 0 & 0 & 0 & 0 & 0 & 1-\varepsilon & \varepsilon/2 & \varepsilon/2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1-\varepsilon & \varepsilon \\ 1-\varepsilon & \varepsilon/3 & \varepsilon/3 & \varepsilon/3 & 0 & 0 & 0 & 0 \\ 0 & 1-\varepsilon & \varepsilon/2 & \varepsilon/2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1-\varepsilon & \varepsilon & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Note that if player 2 believes that player 1 will stop the game at the first node ($T = 0$), then player 2 is indifferent about the action to take, i.e. there is no best response. This explains the uniform probability distribution over the actions of player 2 in the first row of each matrix.

References

1. Camerer CF. 2011 The roundtable series in behavioral economics. In *Behavioral game theory: experiments in strategic interaction*. Princeton, NJ: Princeton University Press. See https://books.google.be/books?id=cr_Xg7cRvdcC.
2. McKelvey RD, Palfrey TR. 1992 An experimental study of the centipede game. *Econometrica* **60**, 803. (doi:10.2307/2951567)
3. Rand DG, Peysakhovich A, Kraft-Todd GT, Newman GE, Wurzbacher O, Nowak MA, Greene JD. 2014 Social heuristics shape intuitive cooperation. *Nat. Commun.* **5**, 3677. (doi:10.1038/ncomms4677)
4. Oosterbeek H, Sloof R, van de Kuilen G. 2004 Cultural differences in ultimatum game experiments: evidence from a meta-analysis. *Exp. Econ.* **7**, 171–188. (doi:10.1023/b:exec.0000026978.14316.74)
5. Grosskopf B, Nagel R. 2008 The two-person beauty contest. *Games Econ. Behav.* **62**, 93–99. (doi:10.1016/j.geb.2007.03.004)
6. Camerer CF, Ho TH, Chong JK. 2004 Behavioural game theory: thinking, learning and teaching. In *Advances in understanding strategic behaviour: game theory, experiments and bounded rationality* (ed. S Huck), pp. 120–180. London, UK: Palgrave Macmillan. (doi:10.1057/9780230523371_8)
7. Jensen K, Call J, Tomasello M. 2007 Chimpanzees are rational maximizers in an ultimatum game. *Science* **318**, 107–109. (doi:10.1126/science.1145850)

8. Bornstein BH, Emler AC. 2001 Rationality in medical decision making: a review of the literature on doctors' decision-making biases. *J. Eval. Clin. Pract.* **7**, 97–107. (doi:10.1046/j.1365-2753.2001.00284.x)
9. Montibeller G, von Winterfeldt D. 2015 Cognitive and motivational biases in decision and risk analysis. *Risk Anal.* **35**, 1230–1251. (doi:10.1111/risa.12360)
10. Marshall JAR, Trimmer PC, Houston AI, McNamara JM. 2013 On evolutionary explanations of cognitive biases. *Trends Ecol. Evol.* **28**, 469–473. (doi:10.1016/j.tree.2013.05.013)
11. Trivers R. 2011 *Deceit and self-deception: fooling yourself the better to fool others*. London, UK: Allen Lane.
12. von Hippel W, Trivers R. 2011 The evolution and psychology of self-deception. *Behav. Brain Sci.* **34**, 1–16. (doi:10.1017/s0140525x10001354)
13. Rusch T, Steixner-Kumar S, Doshi P, Spezio M, Gläscher J. 2020 Theory of mind and decision science: towards a typology of tasks and computational models. *Neuropsychologia* **146**, 107488. (doi:10.1016/j.neuropsychologia.2020.107488)
14. McKay RT, Dennett DC. 2009 The evolution of misbelief. *Behav. Brain Sci.* **32**, 493–510. (doi:10.1017/s0140525x09990975)
15. Lenaerts T, Saponara M, Pacheco JM, Santos FC. 2024 Evolution of a theory of mind. *iScience* **27**, 108862. (doi:10.1016/j.isci.2024.108862)
16. Camerer CF, Ho TH, Chong JK. 2004 A cognitive hierarchy model of games. *Q. J. Econ.* **119**, 861–898. (doi:10.1162/0033553041502225)
17. Smith JM, Price GR. 1973 Logic of animal conflict. *Nature* **246**, 15–18. (doi:10.1038/246015a0)
18. Smith J. 1982 *Evolution and the theory of games*. Cambridge, UK: Cambridge University Press.
19. Sigmund K. 2010 *The calculus of selfishness*. Princeton, NJ: Princeton University. (Princeton series in theoretical and computational biology).
20. Fernández Domingos E, Santos FC, Lenaerts T. 2023 EGTtools: evolutionary game dynamics in Python. *iScience* **26**, 106419. (doi:10.1016/j.isci.2023.106419)
21. Stahl DO. 1993 Evolution of *Smartn* players. *Games Econ. Behav.* **5**, 604–617. (doi:10.1006/game.1993.1033)
22. Kawagoe T, Takizawa H. 2012 Level-k analysis of experimental centipede games. *J. Econ. Behav. Organ.* **82**, 548–566. (doi:10.1016/j.jebo.2012.03.010)
23. Nax HH, Newton J. 2022 Deep and shallow thinking in the long run. *Theor. Econ.* **17**, 1501–1527. (doi:10.3982/te4824)
24. Aumann RJ. 1997 Rationality and bounded rationality. *Games Econ. Behav.* **21**, 2–14. (doi:10.1006/game.1997.0585)
25. Rosenthal RW. 1981 Games of perfect information, predatory pricing and the chain-store paradox. *J. Econ. Theory* **25**, 92–100. (doi:10.1016/0022-0531(81)90018-1)
26. Rapoport A, Stein WE, Parco JE, Nicholas TE. 2003 Equilibrium play and adaptive learning in a three-person centipede game. *Games Econ. Behav.* **43**, 239–265. (doi:10.1016/s0899-8256(03)00009-5)
27. Murphy RO, Rapoport A, Parco JE. 2004 Population learning of cooperative behavior in a three-person centipede game. *Ration. Soc.* **16**, 91–120. (doi:10.1177/1043463104039876)
28. Fey M, McKelvey RD, Palfrey TR. 1996 An experimental study of constant-sum centipede games. *Int. J. Game Theory* **25**, 269–287. (doi:10.1007/BF02425258)
29. Palacios-Huerta I, Volij O. 2009 Field centipedes. *Am. Econ. Rev.* **99**, 1619–1635. (doi:10.1257/aer.99.4.1619)
30. Rand DG, Nowak MA. 2012 Evolutionary dynamics in finite populations can explain the full range of cooperative behaviors observed in the centipede game. *J. Theor. Biol.* **300**, 212–221. (doi:10.1016/j.jtbi.2012.01.011)
31. McKelvey RD, Palfrey TR. 1995 Quantal response equilibria for normal form games. *Games Econ. Behav.* **10**, 6–38. (doi:10.1006/game.1995.1023)
32. McKelvey RD, Palfrey TR. 1998 Quantal response equilibria for extensive form games. *Exp. Econ.* **1**, 9–41. (doi:10.1023/A:1009905800005)
33. Béla E. 2022 Altruism and ambiguity in the centipede game (doi:10.31235/osf.io/93p8s)
34. Krockow EM, Colman AM, Pulford BD. 2016 Exploring cooperation and competition in the centipede game through verbal protocol analysis. *Eur. J. Soc. Psychol.* **46**, 746–761. (doi:10.1002/ejsp.2226)
35. Gamba A, Regner T. 2015 Preferences-dependent learning in the centipede game: the persistence of mistrust. *Eur. Econ. Rev.* **120**, 103316. (doi:10.1016/j.eurocorev.2019.103316)
36. Fawcett TW, Fallenstein B, Higginson AD, Houston AI, Mallpress DEW, Trimmer PC, McNamara JM. 2014 The evolution of decision rules in complex environments. *Trends Cogn. Sci.* **18**, 153–161. (doi:10.1016/j.tics.2013.12.012)
37. Tversky A, Kahneman D. 1982 Judgment under uncertainty: heuristics and biases. *Science* **185**, 1124–1131. (doi:10.1126/science.185.4157.1124)
38. de Weerd H, Verbrugge R, Verheij B. 2013 How much does it help to know what she knows you know? An agent-based simulation study. *Artif. Intell.* **199–200**, 67–92. (doi:10.1016/j.artint.2013.05.004)
39. Devaine M, Hollard G, Daunizeau J. 2014 Theory of mind: did evolution fool us? *PLoS One* **9**, e87619. (doi:10.1371/journal.pone.0087619)
40. Basu K. 1994 The traveler's dilemma: paradoxes of rationality in game theory. *Am. Econ. Rev.* **84**, 391–395. <https://www.jstor.org/stable/2117865>
41. Goeree JK, Holt CA. 1999 Stochastic game theory: for playing games, not just for doing theory. *Proc. Natl Acad. Sci. USA* **96**, 10564–10567. (doi:10.1073/pnas.96.19.10564)
42. Goeree JK, Holt CA. 2004 A model of noisy introspection. *Games Econ. Behav.* **46**, 365–382. (doi:10.1016/s0899-8256(03)00145-3)
43. Stewart AJ, Plotkin JB. 2013 From extortion to generosity, evolution in the Iterated Prisoner's Dilemma. *Proc. Natl Acad. Sci. USA* **110**, 15348–15353. (doi:10.1073/pnas.1306246110)
44. Auger P, de la Parra RB, Sánchez E. 1998 Hawk-dove game and competition dynamics. *Math. Comput. Model.* **27**, 89–98. (doi:10.1016/s0895-7177(98)00009-0)
45. Schuster P, Sigmund K. 1983 Replicator dynamics. *J. Theor. Biol.* **100**, 533–538. (doi:10.1016/0022-5193(83)90445-9)
46. Cressman R, Tao Y. 2014 The replicator equation and other game dynamics. *Proc. Natl Acad. Sci. USA* **111**, 10810–10817. (doi:10.1073/pnas.1400823111)
47. Nowak MA. 2006 Five rules for the evolution of cooperation. *Science* **314**, 1560–1563. (doi:10.1126/science.1133755)
48. Taylor C, Nowak MA. 2007 Transforming the dilemma. *Evolution* **61**, 2281–2292. (doi:10.1111/j.1558-5646.2007.00196.x)
49. Heller Y. 2015 Three steps ahead. *Theor. Econ.* **10**, 203–241. (doi:10.3982/te1660)
50. Crawford VP. 2003 Lying for strategic advantage: rational and boundedly rational misrepresentation of intentions. *Am. Econ. Rev.* **93**, 133–149. (doi:10.1257/000282803321455197)
51. Aue T, Nusbaum HC, Cacioppo JT. 2012 Neural correlates of wishful thinking. *Soc. Cogn. Affect. Neurosci.* **7**, 991–1000. (doi:10.1093/scan/nsr081)
52. Neumann R. 2014 A bounded rationality account of wishful thinking. In *Proc. of the Annual Meeting of the Cognitive Science Society*. vol. **36**. <https://escholarship.org/uc/item/3s74k58c>.
53. Yildiz M. 2007 Wishful thinking in strategic environments. *Rev. Econ. Stud.* **74**, 319–344. (doi:10.48550/arXiv.2011.13846)
54. Vosgerau J. 2010 How prevalent is wishful thinking? Misattribution of arousal causes optimism and pessimism in subjective probabilities. *J. Exp. Psychol.* **139**, 32–48. (doi:10.1037/a0018144)
55. Sanna LJ. 1998 Defensive pessimism and optimism: the bitter-sweet influence of mood on performance and prefactual and counterfactual thinking. *Cogn. Emot.* **12**, 635–665. (doi:10.1080/026999398379484)
56. Bateson M, Desire S, Gartside SE, Wright GA. 2011 Agitated honeybees exhibit pessimistic cognitive biases. *Curr. Biol.* **21**, 1070–1073. (doi:10.1016/j.cub.2011.05.017)
57. Krizan Z, Windschitl PD. 2009 Wishful thinking about the future: does desire impact optimism? *Soc. Personal. Psychol. Compass* **3**, 227–243. (doi:10.1111/j.1751-9004.2009.00169.x)
58. Krizan Z, Windschitl PD. 2007 The influence of outcome desirability on optimism. *Psychol. Bull.* **133**, 95–121. (doi:10.1037/0033-2909.133.1.95)
59. Dekel E, Scotchmer S. 1999 On the evolution of attitudes towards risk in winner-take-all games. *J. Econ. Theory* **87**, 125–143. (doi:10.1006/jeth.1999.2537)
60. Brocas I, Carrillo JD. 2025 Why do children pass in the centipede game? Cognitive limitations v. risk calculations. *Games Econ. Behav.* **150**, 295–311. (doi:10.1016/j.geb.2025.01.003)
61. Johnson DDP, Fowler JH. 2011 The evolution of overconfidence. *Nature* **477**, 317–320. (doi:10.1038/nature10384)

62. Fatima S, Jennings NR, Wooldridge M. 2024 Learning to resolve social dilemmas: a survey. *J. Artif. Intell. Res.* **79**, 895–969. (doi:[10.1613/jair.1.15167](https://doi.org/10.1613/jair.1.15167))
63. Fudenberg D, Imhof LA. 2006 Imitation processes with small mutations. *J. Econ. Theory* **131**, 251–262. (doi:[10.1016/j.jet.2005.04.006](https://doi.org/10.1016/j.jet.2005.04.006))
64. Traulsen A, Nowak MA, Pacheco JM. 2006 Stochastic dynamics of invasion and fixation. *Phys. Rev. E* **74**, 011909. (doi:[10.1103/PhysRevE.74.011909](https://doi.org/10.1103/PhysRevE.74.011909))
65. Hindersin L, Wu B, Traulsen A, García J. 2019 Computation and simulation of evolutionary game dynamics in finite populations. *Sci. Rep.* **9**, 6946. (doi:[10.1038/s41598-019-43102-z](https://doi.org/10.1038/s41598-019-43102-z))
66. Saponara M. 2025 MarcoSaponara/centipede-bias: first public release (doi:[10.5281/zenodo.15838520](https://doi.org/10.5281/zenodo.15838520))
67. Yamagishi T, Yamagishi M. 1994 Trust and commitment in the United States and Japan. *Motiv. Emot.* **18**, 129–166. (doi:[10.1007/BF02249397](https://doi.org/10.1007/BF02249397))
68. Krockow EM, Takezawa M, Pulford BD, Colman AM, Smithers S, Kita T, Nakawake Y. 2018 Commitment-enhancing tools in *centipede* games: evidencing European–Japanese differences in trust and cooperation. *Judgm. Decis. Mak.* **13**, 61–72. (doi:[10.1017/s1930297500008822](https://doi.org/10.1017/s1930297500008822))
69. Krockow EM, Colman AM, Pulford BD. 2016 Cooperation in repeated interactions: a systematic review of Centipede game experiments, 1992–2016. *Eur. Rev. Soc. Psychol.* **27**, 231–282. (doi:[10.1080/10463283.2016.1249640](https://doi.org/10.1080/10463283.2016.1249640))
70. Krockow EM, Takezawa M, Pulford BD, Colman AM, Kita T. 2017 Cooperation and trust in Japanese and British samples: evidence from incomplete information games. *Int. Perspect. Psychol.* **6**, 227–245. (doi:[10.1037/ipp0000074](https://doi.org/10.1037/ipp0000074))